

Figure 2.5: The horizontal and vertical structure of the HadCM3 climate model.

Typically, a global climate model breaks up the surface of the earth into a number of latitude/longitude grid boxes. It divides the atmosphere into layers, from the surface to the stratosphere, and does the same for the ocean, from the surface to the deepest waters (Figure 2.5). At each of the points on this three-dimensional grid in the atmosphere a number of equations, derived from the basic laws of physics, are solved which describe the large-scale evolution of momentum, heat and moisture. Similar equations, but including different variables, are solved for the ocean. The third Met Office coupled ocean-atmosphere GCM, HadCM3, has a resolution over land areas of 2.5° latitude \times 3.75° longitude, with 19 vertical levels in the atmosphere and four layers in the soil. The ocean model has 20 vertical levels and a grid size of 1.25° latitude \times 1.25° longitude. In all, there are about a million grid points in the model. At each of these grid points, equations are solved every time the model steps forward (typically 30 min of *model time*) throughout an experiment which typically lasts 250 *model yr*.

The large ensemble of experiments which form the basis of the UKCP probability projections described in Section 2.3.1 use the *slab model* configuration of HadCM3, known as HadSM3. This represents only the top 50 m of the ocean as one layer and prescribes the effects of ocean heat transport rather than simulating ocean currents explicitly. Hence it is much faster to run on a given computer and so we can run many more experiments. These experiments simulate the long-term *equilibrium* climate (a) at current greenhouse gas concentrations and (b) in a world where these are assumed to be double the current concentrations. Although these simulations do not account for possible changes in ocean circulation, surface and atmospheric processes are widely acknowledged to be the leading drivers of the major features of global patterns of climate change, so slab models are used to provide credible realisations of these patterns. In UKCP09 we are able to run many more experiments (that is, bigger ensembles) using the slab model, and hence explore uncertainties in surface and atmospheric

processes more comprehensively. A smaller ensemble of simulations of time-dependent climate change was also produced with the coupled full-ocean model (HadCM3). Relationships between the change patterns simulated between corresponding variants of the slab model and the full ocean model are then used to *timescale* the slab model results, that is, to convert them into a large ensemble of projections of time-dependent changes from 1951 to 2099, whilst also accounting for uncertainties in the projected geographical patterns due to timescaling. We use additional ensembles of HadCM3 simulations to sample uncertainties in ocean transport, sulphur cycle and land carbon cycle processes, and hence also include the effects of these in the projections. We will return to this topic later in this box, and Chapter 3 discusses it in detail.

Parametrisations in climate models

Many of the most important processes in the climate system (for example the drag exerted by hills as air flows over them, and the formation of clouds) take place at a scale much smaller than the grid size of GCMs — these are called subgrid-scale processes. These cannot therefore be described explicitly, so we develop relationships, known as parametrisations, which estimate them from grid scale variables such as winds, temperature, humidity, etc. which are explicitly described in the model.

We illustrate this by taking the example of cloud amount. This is defined as the proportion of each model grid square which is covered by cloud at each level in the atmosphere. To calculate cloud amount in HadCM3, we use the model's calculated mean temperature and water vapour content for that square and level; this is known as *parametrising* cloud amount in terms of the large scale model variables. Now the equation relating water vapour and temperature to cloud amount contains some parameters, the values of which are based on results from, for example, aircraft measurements or high resolution process models such as cloud resolving models. The values of these parameters are uncertain, and this is a major cause of model uncertainty. So, to quantify this model uncertainty, we vary these parameter values between plausible limits to form variants of a number of configurations of the model, in order to generate the ensembles of simulations which form the primary basis for the PDFs in UKCP09.

But the parametrisation which predicts cloud amount from the modelled large scale variables may be different in models from other centres; not just the parameter values but the actual form of the parametrisation scheme itself; this is illustrated schematically in Figure 2.6. This is an example of a structural difference between models; the effect of structural differences cannot be taken account of using variants of a single model alone. In UKCP09 it is taken into account in the probabilistic projections by using a number of models from other centres, as explained in Chapter 3.

Feedbacks

Basic greenhouse theory tells us that when the concentration of a greenhouse gas, such as CO₂, increases in the atmosphere, it alters the balance between the amount of incoming energy from the sun and that leaving the earth as infrared energy (the radiative balance). Given enough

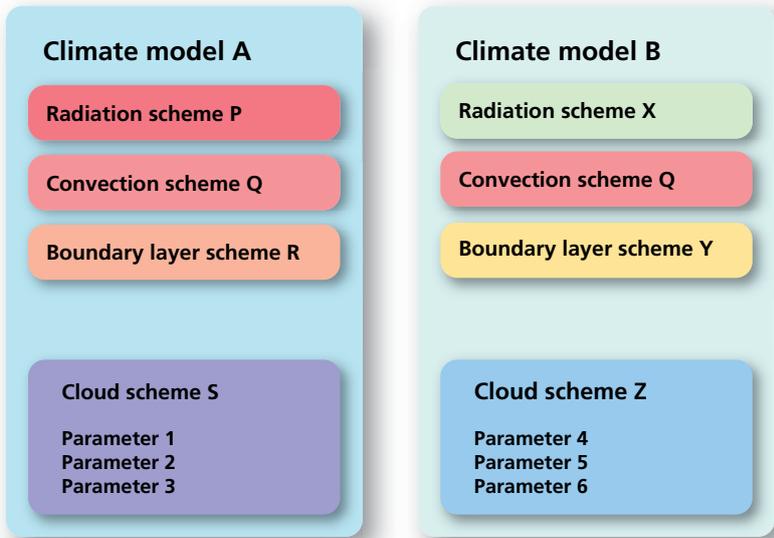


Figure 2.6: Schematic illustration of parametrization schemes in two different climate models, and the parameter values within one scheme (that for cloud). Note that different models may share one or more parametrization schemes; in the diagram this is denoted by the convection scheme.

time, the climate system adjusts to this new condition by increasing the surface temperature of the earth. The direct radiative effect of a doubling the concentration of CO₂ in the atmosphere would eventually cause the surface temperature of the earth to increase by about 1°C. However, once a greenhouse warming starts, a number of consequent changes start to happen which can act to either reduce or increase the direct greenhouse warming; these are known as negative or positive feedbacks respectively.

We illustrate this with some examples. Firstly, as the atmosphere starts to warm due to the direct greenhouse effect, it can “hold” more water vapour — and models indicate that water vapour concentration increases to maintain time-averaged relative humidity (which also depends on temperature) approximately constant as climate change proceeds. As water vapour is a powerful greenhouse gas this effect will further increase warming — a positive feedback. Secondly, as the oceans start to warm some sea-ice will melt. Sea-ice reflects back a lot of solar radiation, but the open ocean it exposes when it melts absorbs more radiation; this will reinforce the original warming effect — another positive feedback. Thirdly, one of the most critical feedbacks, but also one of the most complex, is that due to changes in clouds. In the present climate, clouds have a large effect on climate; high clouds act to increase surface temperatures but low clouds tend to cool climate; the net effect is a cooling one. Greenhouse gas — driven climate change can alter many characteristics of clouds at all levels — their amount and altitude, and the properties of their constituent water droplets and ice crystals, for example. Such changes can alter the radiative properties of clouds — the effect they have on incoming solar radiation and outgoing long wave radiation — and the net effect could be either positive or negative. The last example is that of changes of land surface vegetation (from forests to grassland, for example, or desertification) due to changes in rainfall or temperature which in turn can alter local and global climate. There are many other feedbacks, both positive and negative, in different parts of the climate system.

Feedbacks naturally arise in the climate model because the processes which lead to them (in the second example above this is the formation of sea-ice and its reflectivity) are explicitly represented or parametrised. Many feedbacks take place at a small scale and capturing their overall effect in the model therefore depends upon the parametrisations of small scale processes. Hence the strength of the feedbacks, and thus future changes in climate, will depend on the form of the parametrisation used (part of the model structure), and the values of its constituent parameters. This is one of the main causes of the differences between projections from different models. The methodology developed for the UKCP09 projections is designed to sample these uncertainties, to the extent that this is presently possible, in a systematic way.

Biogeochemical cycles

The carbon cycle and the sulphur cycle represent two important processes in climate change, yet, as with standard processes in the atmosphere and oceans, they carry their own large uncertainties. Here we give an overview of the processes, the uncertainties, and how UKCP09 includes them in the final probabilistic projections; more detail resides in Chapter 3.

The carbon cycle

Currently about half of the emissions of CO₂ from human activities (fossil fuel combustion and land use change) are taken up by sinks on land (vegetation and soils) and in the ocean (seawater and ecosystems within it), leaving the remainder of the CO₂ in the atmosphere where it increases concentrations. But as climate starts to change, carbon sinks can also change, so may be able to absorb more, or less, CO₂ from the atmosphere. For example, as soils warm they increase their respiration of CO₂ back to the atmosphere and their ability to remove CO₂ will weaken, leading to atmospheric concentrations being higher than they would otherwise be — a positive feedback. On the other hand, a warmer climate will encourage the growth of boreal forests which would take up more CO₂ from the atmosphere — a negative feedback. There are a host of such feedbacks, both positive and negative, although the net effect is a positive one. Uncertainties in estimating atmospheric concentrations resulting from emissions were not dealt with in the IPCC Third Assessment Report (TAR) in 2001, and hence could not be taken into account in UKCIP02. In UKCP09 these feedbacks are included, and the uncertainty they add to climate change projections is estimated using two sources of information. Firstly, using variants of the Met Office coupled climate — carbon cycle model with different values for the land carbon cycle parameters within it. Secondly, using results from a project (known as C4MIP) which compared results from a number of international models which include the carbon cycle. Further detail is given in Chapter 3. Note that, although UKCP09 projections include the feedback from both land- and ocean-carbon cycle projections, they only include the effect of uncertainties in the feedback from land, which has been estimated (in C4MIP, see Friedlingstein *et al.* 2006) to be several times greater than that from the ocean component. Because the processes involved in climate — carbon cycle feedback are less well understood, and projections are less constrained by observations, our ability to assess the uncertainty in these is more limited than for other aspects of the climate system.

The sulphur cycle

Sulphur gases emitted from fossil fuel burning, and naturally from the oceans, takes part in chemical reactions in the atmosphere to form small particles — sulphate aerosol. These are eventually removed from the atmosphere by rain and clouds, having a typical lifetime of a few days, but whilst in the atmosphere they can have a substantial cooling effect on climate in a direct and an indirect way. The direct cooling effect arises when a suspension of aerosols in the clear atmosphere reflects back some of the incoming solar radiation before it has a chance to warm the ground. The indirect effect arises from the ability of sulphate particles to act as additional nuclei on which water vapour condenses to form clouds. Such clouds would therefore have more water droplets, each of which (for a given amount of available water) would be smaller — the total surface area would therefore be greater and the cloud would reflect back more solar radiation — a further cooling effect. Both the direct and indirect effects described above are included in the HadCM3 model.

A second indirect effect occurs within sulphate-laden clouds. Because their droplets are smaller than those in clean air, the processes which lead the droplets to grow heavy enough to form rain are slower, and hence the clouds persist (and reflect back solar radiation) longer — a further indirect cooling effect. This is a much more complex process, and is only now becoming understood well enough to be included in models (such as the Met Office earth system model, HadGEM1) but is not included in UKCP09. Because atmospheric sulphate burdens are expected to decline in the future, the omission of this effect may lead to an underestimate of changes in the first few decades of the UKCP09 projections.

Constituents included, and not included, in the probabilistic projections

The atmospheric constituents included in HadCM3, its corresponding simple-ocean configuration and the regional climate model, are shown in Table 2.1. With the exception of the cloud persistence effect of sulphate aerosols, the projected combined effect by 2100 of changes in those constituents not included is unlikely to add a significant amount to overall uncertainty. Similarly, although the Met Office model includes the effect of chemical reactions in the atmosphere which determine concentrations of methane and tropospheric (low altitude) ozone, no attempt was made to estimate the consequent uncertainty in concentrations; this would also be expected to have a minor effect. Uncertainty in the climate effect of northern hemisphere stratospheric ozone changes is also likely to be small relative to those quantified.

In contrast, other components of the methane cycle, such as climate-induced emissions from wetlands, melting permafrost and methane hydrates, do have the potential to modify future climate change significantly. However, these feedbacks are so poorly understood as to make estimates of their effect very uncertain, and hence they are not currently integrated into any climate model.

Constituent	Whether included
Carbon dioxide	Yes
Methane	Yes
Nitrous oxide	Yes
CFCs, PFCs, HFCs, HCFCs, SF6	Major ones
Tropospheric ozone	Yes
Stratospheric ozone	Yes
Sulphate aerosols — direct effect	Yes
Sulphate aerosols — cloud albedo effect	Yes
Sulphate aerosols — cloud persistence effect	No
Black carbon aerosol	No
Organic carbon aerosol	No
Mineral dust	No
Sea salt aerosol	No
Land cover (albedo effect)	No

Table 2.1: The atmospheric constituents included in the Met Office models used for UKCP09.

At a local scale, differences between projections are even more obvious. Figure 2.7 shows, as an example, projections of changes in summer precipitation over the UK from 12 climate models, for the same future time period and same emissions scenario. Rainfall over London shows a reduction of about 60% in the projection from one model, but a small increase in another. Note that, because Figure 2.7 shows only single projections — all that is available from most climate models — natural internal variability contributes to the differences between them.

A similar illustration of model differences was shown in UKCIP02. The differences now are no smaller than those shown 7 yr ago — in other words, there has been no apparent convergence of model projections, despite improvements in climate process representations in models made during this period. For this reason, we cannot assume that continuing model improvements will quickly lead to a narrowing of uncertainty in projections.



Figure 2.7: Changes (%) in summer (June–August) precipitation by the period 2071–2100 compared to 1961–1990, from 12 climate models, each of which took part in the IPCC AR4, all driven with the same SRES A2 emissions scenario.

Planners and decision-makers could, of course, use the range of projections such as those in Figure 2.7 as an estimate of the uncertainty which should be taken into account, and the UKCIP02 report recommended this course of action. Of more use to planners would be some indication of the relative credibility of each of the models, but systematic techniques for doing this are difficult to apply to such a small and diverse set of climate models. In UKCP09 we quantify the uncertainties in projections, giving information on the relative likelihood of different climate change outcomes, in the form of probabilistic projections. In this way, rather than give users a single projection of unknown likelihood, we can show the uncertainty in projections in the form of a probability distribution function or PDF. This shows us the relative probability of temperatures changes of, say 2°C or 3°C at a particular location by a certain time period. The interpretation of this probability is important and is discussed in Box 1.3 and Section 2.5. More usefully, it can be used to estimate the probability of a change being greater or less than some threshold. The method gives probabilities of changes in number of variables, both monthly means and some extremes. PDFs, and an alternative method of presenting the same information, the Cumulative Distribution Function (CDF), are explained in more detail in Box 1.3.

The requirement for probabilistic projections has been recognised by the climate modelling community for some time, and they have begun to develop methods based on projections that are available from a number of climate models – the so called *ensemble of opportunity* (Giorgi and Mearns, 2003; Dessai *et al.* 2005; Goodess *et al.* 2007; CSIRO and Bureau of Meteorology, 2007; Frei, 2007). However, whilst such an ensemble (as in Figure 2.7) is sufficient to demonstrate the requirement for probabilistic projections, it is not sufficient to fulfil it. This is because it is assembled on an ad-hoc basis, and has not been designed to sample modelling uncertainties in a systematic and comprehensive manner. The ensemble of opportunity in Figure 2.7 shows some range of projections, but does not indicate in which part of the range the outcome is likely to lie — it may even be outside the model range. We therefore base the UKCP09 on an alternative approach, which nevertheless uses the information from an international set of climate models, described in outline below and in more detail in Chapter 3.

2.3.1 Accounting for modelling uncertainty in UKCP09

As summarised earlier, uncertainties in model projections arise from an incomplete understanding of processes in the Earth's climate system, and an inadequate representation of these processes in climate models. These representations may be limited not only by physical knowledge but also by, for example, computing resources, and these lead to errors in models, which in turn lead to errors in projections. For convenience we group all these under the heading *modelling uncertainty*.

In UKCP09 we sample uncertainties in a range of processes in the atmosphere and at the surface, the carbon and sulphur cycles, and in the ocean. However, we recognise that uncertainties in atmospheric processes are likely to be the major contributor to overall uncertainty at a local level, and hence these are treated in the greatest detail in the UKCP09 methodology. The development of new techniques to sample atmospheric model errors, and hence account for their effects in driving uncertainty in future projections of climate, is a key aspect of the research underpinning UKCP09. In order to understand the approach, it is convenient to separate sources of model error into two types: structural error and parameter error. The UKCP09 approach seeks to sample uncertainties arising from both of these. In the first case, when building a model the modeller will make choices about its basic structure, such as the grid on which atmospheric

or oceanic motions are resolved, the numerical integration scheme, the set of physical processes included, etc. Many important processes (such as those in clouds) occur on spatial scales too small to be resolved explicitly on the model grid, and therefore have to be represented in models using relationships with large scale variables which are resolved — so-called *sub-grid scale parametrisations*. The nature of the equations used for a given representation is an important component of its structure. Models containing different structural choices will possess different biases in their simulations of climate processes, and hence give different projections of change — this is the structural component of model error. In the second case, having chosen a particular parametrisation scheme to represent a given small scale process, the modeller has then to choose the values of parameters which control how the process operates in that scheme. These parameters are based on a mixture of theory, observations and experimentation, but the available information is seldom precise enough to allow the appropriate value of a given parameter to be accurately known — this gives rise to the parameter component of model error. This is discussed in rather more detail in Box 2.1.

We explore the effects of uncertainties in atmospheric and land model parameters controlling surface and atmospheric processes using one climate model – in this case the Met Office model HadSM3 (a configuration of HadCM3* having a simplified ocean, see Box 2.1). This is done by identifying parameters controlling the detailed processes likely to have the most effect on model projections. Several parameters are selected from each of the schemes in the model's atmosphere and land: layer cloud, convection, radiation, atmospheric dynamics, boundary layer, land surface and sea-ice. This covers uncertainties in the major aspects of the model's physics. Next we ask experts to define a range of plausible values, together with an intermediate estimate, for each of the uncertain parameters.

We then construct a large number (ensemble) of variants of the model, known as a *perturbed physics ensemble*, each of which contains a different choice of parameter values within these expert-specified bounds, and make a projection of climate change with each. As a first step, we can simply take this projection, for a particular quantity such as change in summer rainfall over some location, from each of the ensemble members and present these in the form of a distribution showing how frequently different outcomes occur — this is represented by the blue histogram in Figure 2.8.

In principle, we would build a different model variant with each possible combination of parameter values, but to make climate simulations with each of these variants would require an unfeasibly large amount of computing resources. Hence we chose a manageable number (280) of variants, to cover as comprehensive a range of outcomes as possible. However, the shape of the histogram in Figure 2.8 depends upon which combinations of parameter changes we choose. To predict the response for all the model variants that it was not possible to run, we build an *emulator* of model output, relating it statistically to the model parameters. This is trained on the model results we do have, and then used to estimate values of model output variables that would be obtained for any desired combination of parameter values. The distribution of projections

* HadCM3, the model used as the basis of the UKCP09 projections, was also used for the UKCIP02 scenarios. It might be thought that, six years on, a better model might have been used. However, a recent comparison of climate models with observations (Reichler and Kim, 2008) shows that HadCM3 ranked second out of 17 models compared in CMIP-2 in 2002, but still ranked joint second out of 21 models compared in the CMIP-3 comparison in 2007, where models were compared with a pre-industrial control climate. The most recent Met Office Hadley Centre model does compare somewhat better with observations, but its higher resolution would have drastically reduced the number of ensemble members which could have been run, and hence given a less-comprehensive estimate of uncertainty.

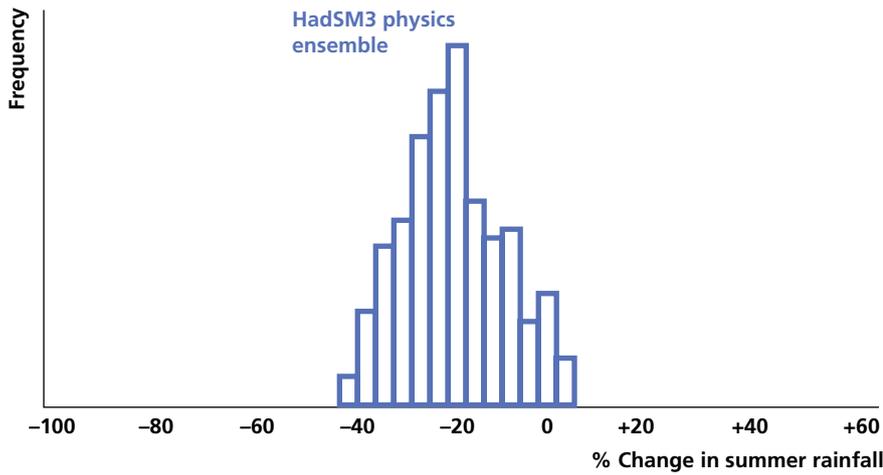


Figure 2.8: Hypothetical histogram showing the frequency of occurrence of different changes in summer rainfall from the 280-member perturbed physics ensemble of HadSM3.

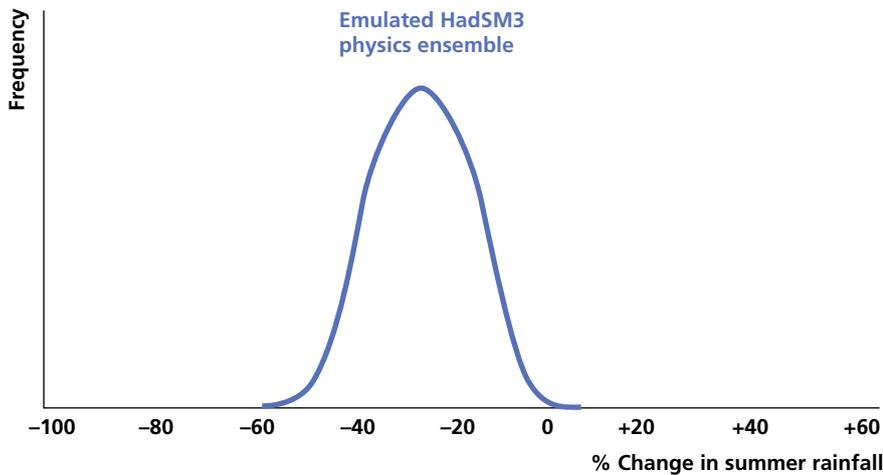


Figure 2.9: Hypothetical distribution showing the frequency of occurrence of different changes from the emulator.

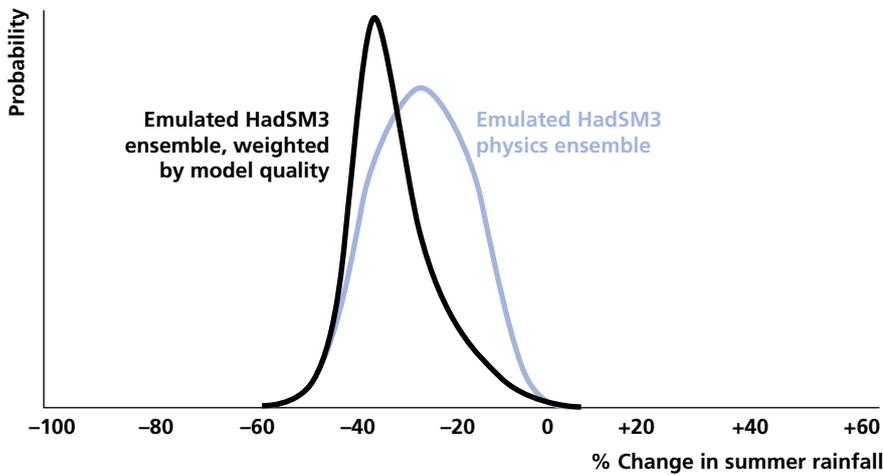


Figure 2.10: Hypothetical distribution showing the probability of different changes from the emulator, weighted according to model credibility based on observations (black curve).

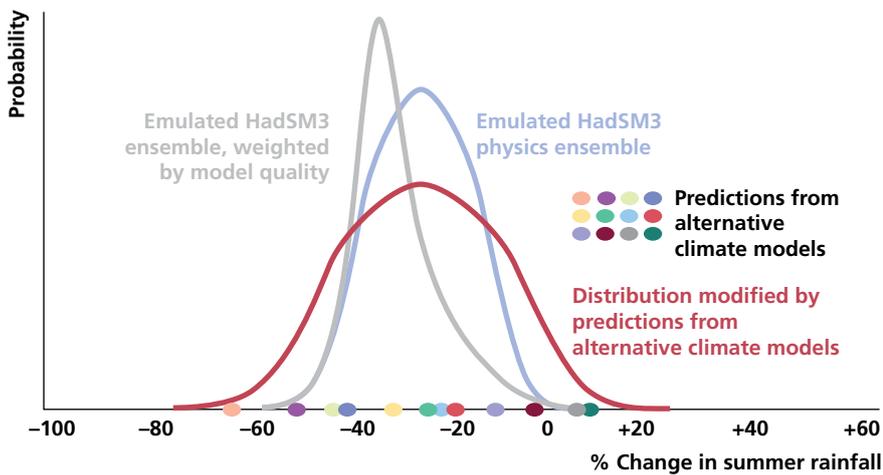


Figure 2.11: The hypothetical probability distribution function of change of summer rainfall (red curve), including projections from both the Met Office perturbed physics ensemble and from alternative international climate models.

from this is illustrated schematically by the blue curve in Figure 2.9, which can take a somewhat different shape from the histogram in Figure 2.8 because the former explores different combinations of parameter values.

Now the model variants will not all give rise to climate simulations of equal credibility, and hence their projections should not be given the same weight. We compare each model's simulation of a wide range of variables for recent climate against observations, and also how well each *hindcasts* large scale patterns of temperature change over the last 90 yr. We use both these pieces of information to weight the projection from each model; this allows us to generate a weighted distribution of outcomes — the black curve in Figure 2.10.*

So far, however, we have described how we use variants of one model to explore the effects of uncertainties in model parameters. However the presence of structural model biases, which cannot be resolved by varying parameters, gives an additional source of uncertainty in model simulations of both past and future climate. This affects both the weights to be assigned to different Met Office model variants, and the spread of possible future projections. We estimate the uncertainty due to these structural errors by using our perturbed physics ensemble to *predict* the results of an alternative set of twelve climate models (all of which have participated in intercomparison exercises such as IPCC AR4) which contain structural assumptions partly independent of those made in the Met Office model. Projections from each of these alternative models are indicated schematically by the coloured dots on Figure 2.11; note that each alternative model is represented by a single projection as no ensemble projections were available. Following IPCC AR4, we assume each of the alternative models has equal validity, bearing in mind that we could not weight the alternative models by re-using the observations employed in determining weights for Met Office model variants, as such double-counting would risk over-constraining our projections.

We assume that differences between the results of the *nearest* few variants of the Met Office model and each of these alternative models gives a reasonable sample of possible differences between the Met Office model and the real world, and hence modify our future projections to account for the resulting estimate of structural model error. These results are then incorporated into our uncertainty analysis, based on a statistical framework devised by Goldstein and Rougier (2004), discussed in Chapter 3. This allows us to create a probability distribution function accounting for uncertainties arising from both model parameters and structural errors, and constrained by observations, shown as the red curve in Figure 2.11.

The above description is an enormously simplified explanation of the methodology. As mentioned earlier, the large ensemble of about 280 members, described above, can only be run using a model configuration with a simple representation of the ocean (known as a slab model, see Box 2.1) which is suitable for the simulation of the long-term *equilibrium* response to an assumed doubling of carbon dioxide, but not for the simulation of time-dependent climate change. Hence additional time-dependent (that is, continuous from 1950 to 2099) simulations are undertaken using the model configuration with atmosphere coupled to a full dynamical ocean (HadCM3). The results from these experiments are used in a technique for matching equilibrium and time-dependent patterns of change so that the very large ensemble of projections using the slab model can be *timescaled*. Further simulations are also needed to sample uncertainties

* Note that in practice the methodology does not involve creation of an interim weighted distribution (as shown in Figure 2.10), prior to the addition of the effects of structural model error; the discussion is presented this way to emphasise the key inputs to the calculations.

arising from ocean transport, carbon cycle and sulphur cycle processes. Finally, to make the projections suitable for impacts and adaptation assessments, we use a further ensemble of the Met Office regional climate model (HadRM3) to *downscale* the projections from the global Met Office model to a resolution of 25 km. A more detailed description of the full methodology is given in Chapter 3. The methodology involves a number of expert choices (for example, the range of values taken for model parameters, and their distribution), the sensitivity to which needs to be tested to establish the robustness of the results. Examples of such sensitivity tests are given in Annex 2.

The relative size of the various contributing factors to the total uncertainty (and hence to the width of the PDF) will be different for different locations, time periods, type of spatial averaging, etc; this is discussed in Annex 2. Figure 2.12 shows two specific examples of the relative contributions, in the case of changes to mean winter precipitation by the 2080s under the Medium emissions scenario, for 25 km squares in south-west England and the west of Scotland. Here we have combined* the proportions of uncertainties due to model parameter values, model structure, the carbon cycle, aerosol physics and ocean physics, and termed this contribution *model uncertainty*. Natural internal variability (*chaos*) is labelled as natural variability. The remaining slice of the pie arises from the timescaling and downscaling procedures in the methodology described above. As can be seen, in these examples modelling uncertainty dominates the other contributions — although this is not true everywhere. A closer time period (the 2020s) would show a relatively bigger contribution from natural variability, and different choices of variables, locations and emissions scenarios would give different pie chart structures. Note that the uncertainty in emissions is not included; this is handled by giving different probability projections for each of three emissions scenarios as described later in this chapter.

The presentation of information in probabilistic terms, rather than giving users a single projection for a given emissions scenario, is a major change in the nature of climate change projections. Whilst they are undoubtedly more complicated to grasp conceptually, and their application in practice demands more of the user, probabilistic projections are a more honest way of representing the substantial uncertainties that are discussed above. Because it is so important to understand, we repeat here the point made in Chapter 1, that a probability given in UKCP09 is not the same as the probability of a given number arising in a game of chance, such as rolling dice. Instead, it is a measure of the degree to which a particular level of future climate change is consistent with the information (observations and model simulations) used in the analysis, that is, the evidence.

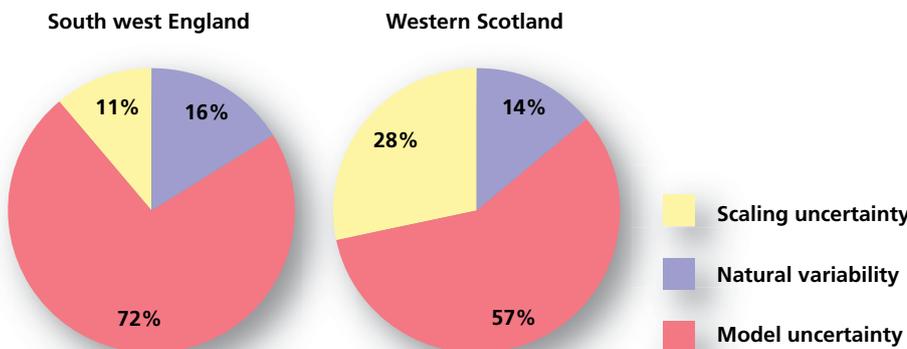


Figure 2.12: The relative contributions to overall uncertainty in change to winter-mean precipitation for 25 km squares in south west England (left) and west Scotland (right) by the 2080s under the Medium emissions scenario, from natural variability, modelling uncertainty and scaling uncertainty. (Contributions do not total 100% due to rounding errors.)

* Because of the way contributions are divided up in Annex 2, this aggregation is a close approximation to, but does not exactly cover, all the terms in model uncertainty.

2.4 Uncertainty due to future emissions

Previous UKCIP reports on climate change projections have discussed uncertainty due to future emissions, and this uncertainty continues to apply to the climate projections in this report. The pathway of future emissions of greenhouse gases (CO₂, methane, nitrous oxide, etc.) and aerosols (or aerosol precursor emissions such as sulphur dioxide) will depend upon many socioeconomic factors such as changes in population, GDP, and energy use, and in technical developments which might influence carbon intensity (the amount CO₂ per unit of energy generated). IPCC published a Special Report on Emissions Scenarios (SRES) (Nakićenović and Swart, 2000), in which climate-relevant emissions were calculated based on a number of *storylines*, each describing a possible pathway of how the world might develop. All scenarios are *non-interventionist*, that is they assume no political action to reduce emissions in order to mitigate climate change; differences between them arise purely from different assumptions about future socioeconomic changes.

There is no agreed method with which to assign a relative probability to different future emissions; SRES made it clear that no relative probability could be attached to different emissions scenarios, but neither were they to be assumed as equally probable (see Annex 1). (Strictly speaking, being scenarios, they have no probability.) This means that the uncertainty due to future emissions cannot be incorporated into a probabilistic projection. However, the uncertainty associated with future emissions is recognised in UKCP09 by giving probabilistic projections which correspond to each of three different emissions scenarios, High, Medium and Low. These scenarios correspond to three of the *marker scenarios* in SRES: A1FI, A1B and B1 respectively, as decided following consultation. This is a change from UKCIP02, where four emissions scenarios were used corresponding to SRES A1FI, A2, B2 and B1. Figure 2.13 shows emissions of CO₂ from the scenarios used in UKCIP02 and UKCP09. Each scenario also includes emissions of other greenhouse gases, and of sulphur dioxide which creates sulphate aerosols that cool climate. Although the three UKCP emissions scenarios span the range of marker scenarios in SRES, there are additional scenarios, both higher and lower, that they do not encompass.

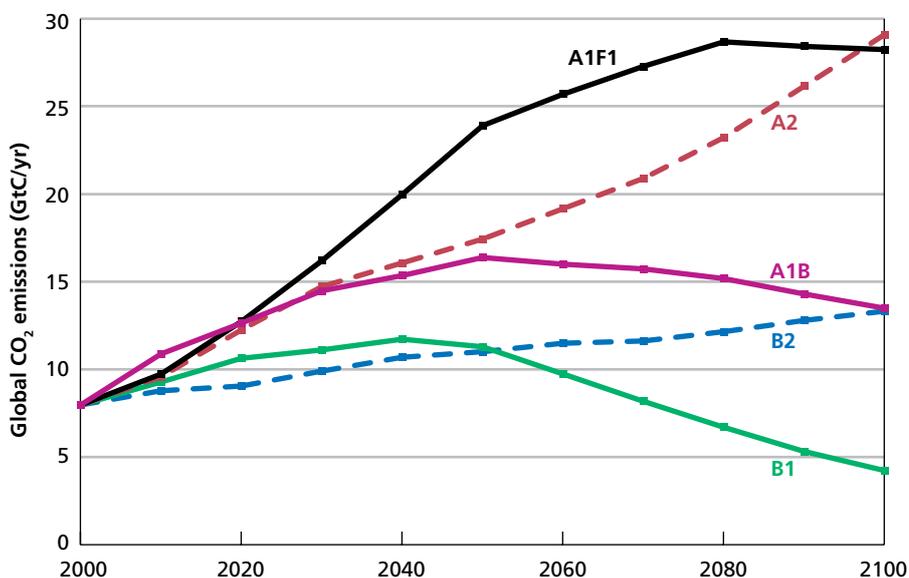


Figure 2.13: Global annual CO₂ emissions (expressed as gigatonnes of carbon) under the three IPCC SRES *marker scenarios* used in UKCP09: A1FI (black: High emissions), A1B (purple: Medium emissions) and B1 (green: Low emissions). Also shown dotted are two SRES emissions scenarios used in UKCIP02 but not in UKCP09: A2 (red: Medium-High Emissions) and B2 (blue: Medium-Low Emissions).

Additional uncertainties arise from the way in which the SRES emissions scenarios were developed, both in the underlying storylines of future changes in society, economies, technology, etc., and in the way in which the emissions are developed from the storylines. These uncertainties are considered here to be part of the overall uncertainty in future emissions.

More detail on the three SRES emissions scenarios, and the socioeconomic futures which underlie them, is given in Annex 1. Of course the question of how to handle results from the three projections from the different emissions scenarios in a risk assessment still remains an issue for users, and this is discussed in the User Guidance.

The differences in projections of global temperature over land which arises from different future emissions is illustrated in Figure 2.14, using the average of 17 variants of the HadCM3 model. Not surprisingly, the High emissions scenario results in the greatest warming by 2100, and the Low emissions scenario gives the smallest warming. But also evident is the relative insensitivity of warming to emissions scenario, over the period to about 2040. This is partly due to the smoothing effect of the long effective lifetime of CO₂ and the thermal inertia of the climate system, but also partly due to the offsetting effects of warming

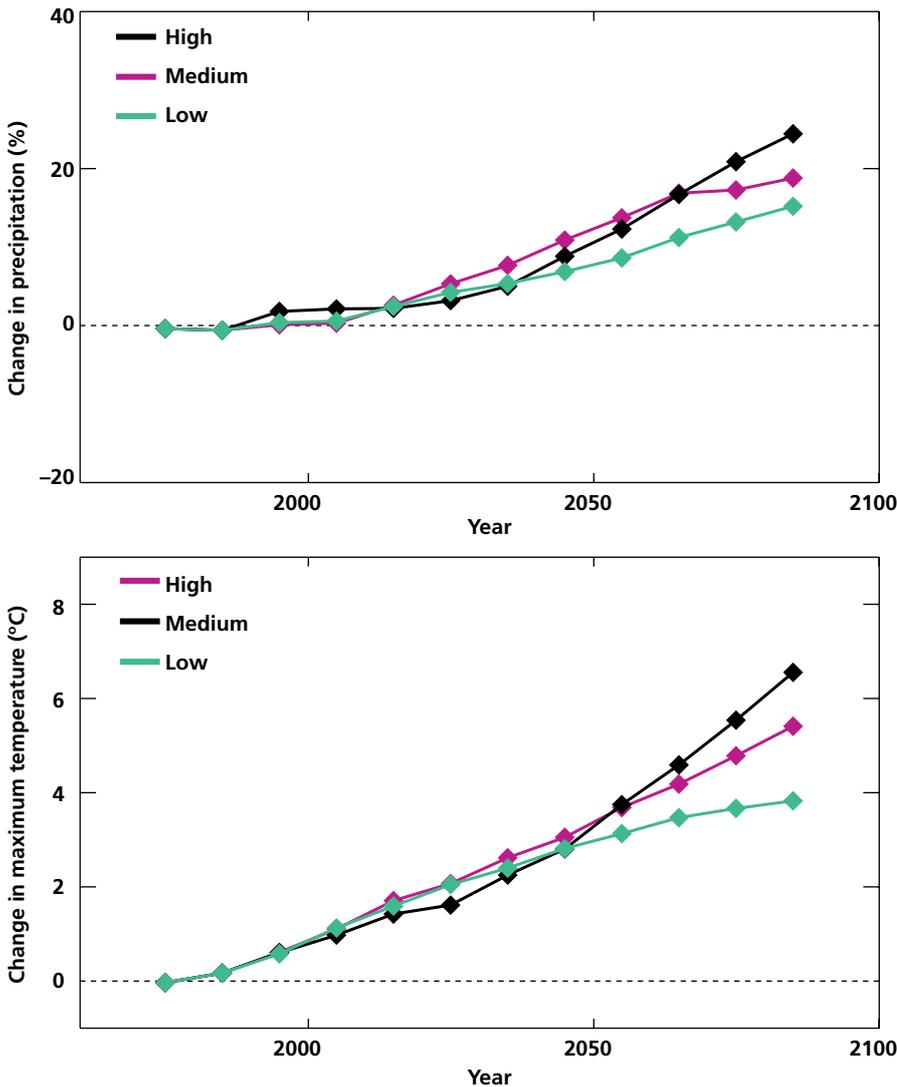


Figure 2.14 Changes in winter-mean precipitation (upper panel), and summer mean daily maximum temperature (lower panel) over Wales, averaged from 17 variants of the HadCM3 global model, for each of three different future emissions scenarios. (Because the purpose here is to show the relative insensitivity of climate change over the next few decades to emissions, the graphs do not reflect the uncertainties in future CO₂ concentrations which are taken into account in UKCP09.)

greenhouse gases and cooling sulphate aerosols in the scenarios. However, after the middle of the century, projections based on the three emissions scenarios become increasingly different.

2.5 Uncertainties in UKCP09 probabilistic projections and future prospects

The procedure used in UKCP09 to convert the ensembles of climate model simulations into probabilistic estimates of future climate necessitates a number of expert choices and assumptions (see Chapter 3 and Annex 2). This implies that the probabilities we specify are themselves uncertain. A system for projecting future climate (unlike one for short-range weather forecasting) cannot be verified on a large sample of past cases. Nevertheless it is possible to check whether or not our probabilistic estimates are robust to reasonable variations within these assumptions; results from some such sensitivity tests are shown in Annex 2.

Although it is important that prospective users understand the limitations and caveats, it is also worth emphasising that (a) current models are capable of simulating many aspects of global and regional climate with considerable skill (see Annex 3); and (b) they do capture, albeit imperfectly, all the major physical and biogeochemical processes known to be likely to exert a significant influence on global and regional climate over the next 100 yr or so.

As explained in the previous section, there are several components of uncertainty which contribute, in varying proportions, to the width of the PDF of change in a particular variable (for a given emissions scenario, location, etc.). These can be thought of as being in three categories:

- uncertainty due to natural variability
- statistical uncertainty inherent in the UKCP09 methodology
- modelling uncertainty (arising from our lack of understanding of the climate system and our inability to model it perfectly) — which includes the carbon cycle, sulphur aerosols and ocean heating.

In the conclusion to Annex 2 we explain how each of these could be reduced in future. By initialising models with recent climate, we should be able to reduce uncertainty due to natural variability, especially for the next 10–20 yr. For long term projections, natural variability represents an irreducible contribution to the overall uncertainty. Uncertainty in the statistical methodology could be reduced with a sizeable increase in computing power. Modelling uncertainty should reduce as our understanding of the climate system and our ability to represent it in climate models gets better, although history shows that this is likely to be slow.

The consequence of these expected improvements is that the shape of a given PDF is likely to change in the future. Users need to understand clearly that, if they choose to adapt to a climate change corresponding to a specific probability level, this is likely to change in future projections — and the changes are likely to be greater at the extremes of probability levels (that is, 10 and 90%). If our understanding of climate processes, and model representations of them, does not change substantially in future, then we foresee a general reduction in uncertainties (except that due to long-term natural variability) because of improvements in our ability to represent processes currently modelled and we would hence expect the shape of the PDF to change, with a reduction in its width. However, we do not know in what way this reduction in width will occur;

in particular it may not be towards what are the most likely values in UKCP09. Although we cannot say what the next generation of PDFs will look like, it is likely that the spread of plausible changes they would indicate would be encompassed by the corresponding PDFs shown in UKCP09. Thus, in the absence of any major change in model projections, users who are incorporating the probabilities given in UKCP09 into their decision making are likely to find that their decisions are robust to changes in the next generation of projections.

On the other hand, there is also the potential for uncertainties to become greater if processes not yet included, or included imperfectly, in the models turn out to exert a substantial influence on climate change. Less than a decade ago, for example, carbon cycle feedbacks were not included in models, yet these are now known not only to change the projections substantially but also to add significantly to the uncertainty in them — which is why they are included in UKCP09. Further such effects, for example, methane feedbacks from land and oceans or the dynamics of ice sheets, may be shown to be important in due course. Uncertainties could also widen if future (improved) models reveal that a process which is represented in the current generation of models, but with a common bias, turns out to exhibit a larger response to man-made forcing than current models suggest (see Box 2.1). However, the consistency between model simulations and observations of change over the last century provides some reassurance that any unknown processes are unlikely to change projections fundamentally, at least for the next few decades.

An obvious follow-up question is: should decisions be made now, based on UKCP09 projections, or should they be delayed in the hope that better projections will be available in a few years time? The risk of deferring a decision is something that can be assessed using the UKCP09 projections. How rapidly will climate projections change in the future? Although modellers have improved many aspects of their models over the past decade or so, the current range of changes over the UK (Figure 2.7) is not significantly narrower than that shown in UKCIP02. In practice, the prospects for better projections will depend on which aspects of future climate users are most interested in. The width of the PDFs in UKCP09 are substantial even for the next few decades, due mainly to natural variability, and grow larger through the century due to uncertainties in climate feedbacks. It may be possible to reduce short-term uncertainties with higher resolution models which may simulate better (for example) the North Atlantic storm track, and by starting model experiments with the recently observed state of the ocean. However, this may not improve projections of (say) changes in surface temperature a hundred years ahead; at these lead times improved projections would come from more faithful representations of climate feedbacks and the carbon cycle in models. Dialogue between decision makers and climate scientists, on the potential for emerging research to update projections, will be essential. However, we reiterate the key point made earlier that the UKCP09 methodology is designed to capture known uncertainties in the climate system built into the current generation of climate models, and is the most comprehensive approach to do so to date. The UKCP09 projections can make a useful contribution to assessing risks posed by future climate; they are appropriate for informing decisions on adaptation to long-term climate change which need to be taken on the basis of current knowledge, and the uncertainty quantified in them is likely to be a conservative estimate.

2.6 References

- CCIRG (1991). *The potential effects of climate change in the United Kingdom*. Department of the Environment, HMSO, London.
- CCIRG (1996). *Review of the potential effects of climate change in the United Kingdom*. Department of the Environment, HMSO, London.
- CSIRO and Bureau of Meteorology (2007). *Climate change in Australia*. Technical Report, 140pp. www.climatechangeinaustralia.gov.au
- DeFries, R. S., Bounoua, L. & Collatz, G. J. (2002). Human modification of the landscape and surface climate in the next fifty years. *Global Change Biology*, **8**, 438–458.
- Dessai, S., Lu, X. & Hulme, M. (2005). Limited sensitivity analysis of regional climate change probabilities for the 21st century. *Journal of Geophysical Research*, **110**, D19108.
- Dessai, S. & Hulme, M. (2008). How do UK climate scenarios compare with recent observations? *Atmospheric Science Letters*, **9**(4), 189–195.
- Frei, C. (2007). Ausschnitt aus einer modellierten Temperaturveränderungskarte der CH2050 — Grundlagenszenarien. MeteoSchweiz, Zurich.
- Friedlingstein, P., Cox, P. M., Betts, R. A., Bopp, L., von Bloh, W., Brovkin, V., Cadule, P., Doney, S., Eby, M., Fung, I., Bala, G., John, J., Jones, C. D., Joos, F., Kato, T., Kawamiya, M., Knorr, W., Lindsay, K., Matthews, H. D., Raddatz, T., Rayner, P., Reick, C., Roeckner, E., Schnitzler, K. G., Schnur, R., Strassmann, K., Weaver, A. J., Yoshikawa, C. & Zeng, N. (2006). Climate-carbon cycle feedback analysis, results from the C4MIP model intercomparison. *Journal of Climate*, **19**, 3337–3353 (doi:10.1175/JCLI3800.1).
- Gedney, N. & Valdes, P. J. (2000). The effect of Amazonian deforestation on the northern hemisphere circulation and climate. *Geophysical Research Letters*, **27**, 3053–3056.
- Giorgi, F. & Mearns, L. O. (2003). Probability of regional climate change based on the reliability ensemble averaging method. *Geophysical Research Letters*, **30**(12), 1629.
- Goldstein, M. & Rougier, J. C. (2004). Probabilistic formulations for transferring inferences from mathematical models to physical systems. *SIAM Journal of Scientific Computing*, **26**, 467–487.
- Goodess, C. M., Hall, J., Best, M., Betts, R., Cabantous, L., Jones, P. D., Kilsby, C. G., Pearman, A. & Wallace, C. J. (2007). Climate scenarios and decision making under uncertainty. *Built Environment*, **33**(1), 10–30.
- Hansen, J., Sato, M., Ruedy, R., Lacis, A., Asamoah, K., Borenstein, S., Brown, E., Cairns, B., Caliri, G., Campbell, M., Curran, B., de Castro, S., Druyan, L., Fox, M., Johnson, C., Lerner, J., McCormick, M. P., Miller, R., Minnis, P., Morrison, A., Pandolfo, L., Ramberran, I., Zaucker, F., Robinson, M., Russell, P., Shah, K., Stone, P., Tegen, I., Thomason, L., Wilder, J. & Wilson, H. (1996). *A Pinatubo climate modeling investigation*, in NATO ASI Series Volume, Subseries I Global Environment Change. Fiocco, G., Fua, D. & Visconti, G. (Eds). Springer-Verlag.
- Hulme, M. & Jenkins, G. J. (1998). *Climate change scenarios for the UK; Scientific Report*. UKCIP Technical Report No. 1, Climatic Research Unit, Norwich, UK, 80 pp.
- Hulme, M., Jenkins, G. J., Lu, X., Turnpenny, J. R., Mitchell, T. D., Jones, R. G., Lowe, J. A., Murphy, J. M., Hassell, D., Boorman, P., McDonald, R. & Hill, S. (2002). *Climate change scenarios for the United Kingdom: The UKCIP02 Scientific Report*. Tyndall Centre for Climate Change Research, School of Environmental Sciences, University of East Anglia, Norwich, UK 120 pp.
- Hulme, M. & Dessai, S. (2008). Negotiating future climates for public policy: a critical assessment of the development of climate scenarios for the UK. *Environmental Science and Policy*, **11**, 54–70.

- IPCC (2007). *Climate Change 2007: The physical science basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Solomon, S., Qin, D., Manning, M., Marquis, M., Averyt, K. B., Tignor, M. & Miller, H. L. (Eds). Cambridge University Press, Cambridge, UK and New York, NY, USA, 996 pp.
- Jenkins, G. J., Perry, M. & Prior, J. (2007). *The climate of the United Kingdom and recent trends*. Met Office, Exeter, UK.
- Keenlyside, N. S., Latif, M., Jungclaude, J., Kornblueh, L. & Roeckner, E. (2008). Advancing decadal-scale climate prediction in the North Atlantic sector. *Nature*, **453**, 84–88.
- Lowe, J. L., Howard, T., Jenkins, G. J., Pardaens, A., Holt, J., Wolf, J., Reeder, T. & Dye, S. (2008). *UK Climate Projections science report: Marine and coastal projections*. Met Office, Exeter, 96 pp.
- Mearns, L. O., Giorgi, F., Whetton, P., Pabon, D., Hulme, M. & Lal, M. (2003). Guidelines for use of climate scenarios developed from regional climate model experiments. IPCC. http://www.ipcc-data.org/guidelines/dgm_no1_v1_10-2003.pdf
- Nakićenović, N. & Swart, R. (Eds.) (2000). *Special Report on Emissions Scenarios*. A Special Report of Working Group III of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, UK and New York, NY, USA, 599 pp.
- Reichler, T. & Kim, J. (2008). How well do coupled models simulate today's climate? *Bulletin of American Meteorological Society*, **89**(3), 303–311.
- Santer, B. D., Mears, C., Wentz, F. J., Taylor, K. E., Gleckler, P. J., Wigley, T. M. L., Barnett, T. P., Boyle, T. S., Bruggemann, W., Gillett, N. P., Klein, S. A., Meehl, G. A., Nozawa, T., Pierce, D. W., Stott, P. A., Washington, W. M. & Wehner, M. F. (2007). Identification of human induced changes in atmospheric moisture content. *Proceedings of the National Academy of Sciences*, **104**, 15248–15253.
- Smith, D. M., Cusack, S., Colman, A. W., Folland, C. K., Harris, G. R. & Murphy, J. M. (2007). Improved surface temperature prediction for the coming decade from a global climate model. *Science*, **317**, 5839, 796–799 (doi:10.1126/science.1139540).
- Soares-Filho, B. S., Nepstad, D. C., Curran, L. M., Cerqueira, G. C., Garcia, R. A., Ramos, C. A., Voll, E., McDonald, A., Lefebvre, P. & Schlesinger, P. (2006). Modelling conservation in the Amazon basin. *Nature*, **440**, 520–523.
- Stainforth, D. A., Allen, M. R., Tredger, E. R. & Smith, L. A. (2007). Confidence, uncertainty and decision-support relevance in climate predictions. *Philosophical Transactions of the Royal Society A*, **365**, 2145–2161.
- Stott, P. A., Jones, G. S. & Mitchell, J. F. B. (2003). Do models underestimate the solar contribution to recent climate change? *Journal of Climate*, **16**(24), 4079–4093.
- West, C. C. & Gawith, M. J. (Eds) (2005). *Measuring progress: preparing for climate change through the UK Climate Impacts Programme*. UKCIP, Oxford. ISBN 0-9544830-5-7.
- Wilby, R. L. (2008). Constructing climate change scenarios of Urban Heat Island intensity and air quality. *Environment and Planning B: Planning and Design*, **35**, 902–919.
- Willows, R. & Connell, R. K. (2003). *Climate adaptation: Risk uncertainty and decision-making*. UKCIP Technical Report, UK Climate Impacts Programme, Oxford, UK, 162 pp.
- Zhang, X., Zwiers, F. W., Hegerl, G. C., Lambert, F. H., Gillett, N. P. & Solomon, S. (2007). Detection of human influence on 20th century precipitation trends. *Nature*, **448**, 461–465.

3 Construction of probabilistic climate projections

The Met Office Hadley Centre has designed a methodology to provide probabilistic projections for UKCP09 which reflect major known uncertainties in relevant climate system processes. The method uses large ensembles of climate model projections, which are processed using advanced statistical methods to generate thousands of plausible climate outcomes, which are then weighted using historical observations.

This chapter provides a comprehensive review of the methodology used to construct the UKCP09 probabilistic projections, for readers requiring a more complete scientific insight into their basis. It is necessarily written assuming a higher level of scientific understanding than other chapters, although it does not seek to document each aspect of the method to the level of technical detail that would appear in a specialist journal paper. Published papers (cited below where relevant) are already available for some components of the method, and will be provided for remaining components in due course. A technical note will also be supplied after the launch of the projections (by October 2009, contingent on the demand for post-launch scientific advice from users), giving a mathematical description of the methodology to supplement the qualitative description given in this chapter.

Section 3.2 describes the elements of the method, and Section 3.3 provides a discussion of the nature, credibility and interpretation of the projections. A short, less technical summary of this material can also be found in Chapter 2, Section 2.2.

3.1 Introduction

It is clear from Chapters 1 and 2 that future climate over the UK (and elsewhere) will be influenced by an array of factors. Some of these affect external forcing of climate through changes to the Earth's radiation balance resulting from natural changes (e.g. volcanic eruptions or variations in solar output) or man-made changes (emissions of greenhouse gases, aerosols and their precursors), while others affect physical and biogeochemical feedback processes which enhance or reduce the response to this forcing. In addition, internal climate variability exerts

a significant influence on climate, in addition to the effects of forced changes. All of these factors introduce uncertainty into projections of future climate because none of them can be predicted perfectly. This is due, in general, to imperfect knowledge of either the detailed behaviour or the current observed states of the relevant systems.

We currently have no agreed method of quantifying the relative likelihood of alternative pathways for future man-made emissions (Section 2.4). For UKCP09, we therefore focus on the task of estimating distributions of future changes in climate for each of three specific emissions scenarios (SRES A1FI, A1B and B1, explained in Section 2.4 and Annex 1, and referred to elsewhere in UKCP09 as High, Medium and Low). These scenarios assume no future changes in natural external forcing, apart from a prescribed repetition of the 11-yr cycle of solar insolation based on past observations. Regional climate changes in response to these emissions will be determined by complex interactions between a number of Earth System processes, plausible projections of which require the use of detailed three-dimensional global climate models (GCMs). As discussed in Section 2.3, ensemble approaches provide an obvious method of exploring the uncertainties associated with GCM projections. Multimodel ensembles (MMEs, e.g. Meehl *et al.* 2005), constructed by pooling projections from alternative GCMs developed at different modelling centres, provide a valuable indication of the range of possible future changes. However, stakeholders faced with climate-sensitive policy and adaptation decisions will typically require more than a simple specification of a possible range (Pittock *et al.* 2001). This is widely recognised in the climate science community, and consequently methods have been suggested to derive probability distributions for regional changes from MME results (e.g. Tebaldi *et al.* 2005; Greene *et al.* 2006; Furrer *et al.* 2007; Watterson, 2008), giving estimates of the relative probability of different future outcomes within the envelope of possible changes. Motivations for such approaches stem from results showing that combining projections from different models can increase the skill of historical climate simulations (e.g. Reichler and Kim, 2008) or seasonal forecasts (e.g. Hagedorn *et al.* 2005), because the errors in different models are partially independent. Furthermore, the models are assembled from a large pool of alternative components, thus sampling to some extent the effects of variations in basic structural assumptions such as choice of model grid, numerical integration scheme or the fundamental physical assumptions employed in the parameterisation of sub-grid scale processes such as convection, boundary layer transports, cloud and precipitation formation, etc. (see Box 2.1). However, multimodel ensembles are rather small in size, consisting typically of 10–20 models, some of which might be run several times from different initial states. Also, the set of models is assembled on an opportunity basis, not being designed to sample systematically some underlying space of possible model formulations (Allen and Stainforth, 2002). This creates the need for substantial assumptions in converting their results into estimated probabilities for climate change, essentially because it is not clear how to identify a distribution of possible outcomes of which the MME is a sample. Different studies address this issue in different ways, and therefore generate significantly different results (see Tebaldi and Knutti, 2007).

Another issue is that probabilistic projections are conditional on the set of uncertainties sampled in the ensemble simulations. In order to provide a credible basis for decision making, a critical prerequisite is that these are designed to sample all sources of uncertainty known to be likely to exert a significant influence on climate over the time frame of interest (here, the 21st century). For a given scenario of future emissions, these would include internal climate variability and uncertainties in atmospheric and oceanic processes, which give rise to different

realisations of 21st century climate in the latest MME produced for the IPCC AR4 (Figure 2.5). However additional sources of uncertainty, notably carbon cycle feedbacks (Box 2.1) and the uncertainty in downscaling GCM simulations to local scales, also need to be considered. In order to produce probabilistic projections for UKCP09, we have therefore developed a new approach aimed at sampling the key uncertainties systematically, using a purpose-built set of ensemble simulations involving several different configurations of the HadCM3 climate model.

The method is based on the notion of the *perturbed physics ensemble* (PPE), in which alternative variants of a single GCM are created by altering the values of uncertain model parameters (Murphy *et al.* 2004; Stainforth *et al.* 2005). These parameters control important small scale processes in the model (such as the formation and precipitation of cloud droplets, the reflectivity of sea ice or the transfer of heat, moisture or momentum between the surface and the atmosphere), and are uncertain because we lack sufficiently detailed observations or sufficiently precise theoretical understanding to constrain their values accurately. A major advantage is that PPEs can be designed to ensure that all the key process uncertainties are sampled in a manner consistent with current scientific understanding. This is achieved by asking experts to identify which model parameters control the key processes, and then to specify distributions for the chosen parameters, consistent with the present state of knowledge concerning the identified processes. We can then construct a set of ensemble runs which select alternative values of the parameters drawn from these distributions, ensuring that the relevant uncertainties are well sampled.

The PPE approach therefore facilitates the construction of probabilistic projections consistent with current understanding of model uncertainties (Section 3.3), and it is also possible to test the sensitivity of the results to reasonable variations in the definition of the *space of possible model variants* implied by the specified distributions for model parameters (see Annex 2). However, the model on which the PPE is based (in our case HadCM3) will inevitably contain some structural errors in its physical representation of the real climate system, which cannot be resolved by varying the model parameters (Murphy *et al.* 2004). These structural errors determine how informative the model simulations are about the real system, so it is critical to account for the additional uncertainty implied by their presence (Goldstein and Rougier, 2004). We address this by using our PPE results to *predict* the results of members of a multimodel ensemble developed at other modelling centres, and containing structural assumptions partially independent of HadCM3. This allows us to estimate the effects of structural errors (subject to assumptions discussed in Section 3.2.8), and to present probabilistic projections which combine information from both perturbed physics and multi-model ensemble results.

The methodology is described in Section 3.2, this being a somewhat abridged (though also updated) version of that given by Murphy *et al.* (2007). Section 3.3 provides a brief summary of key strengths and limitations of our approach, and a discussion of how the probabilistic climate change estimates it provides for UKCP09 should be interpreted by users. The robustness of these estimates to plausible variations in key assumptions is discussed in Annex 2.

3.2 Methodology

3.2.1 Overview

The method is based on a general statistical framework for the derivation of probabilistic projections of real systems from simulations carried out using

complex but imperfect models of those systems (Goldstein and Rougier, 2004; Rougier, 2007). The approach is Bayesian in nature, seeking to estimate the relative credibility of different future outcomes by updating subjective estimates of uncertainty specified before the experiments with evidence from observations. This is achieved by first defining a space of possible variants of the model (through distributions for model parameters consistent with expert knowledge — see Section 3.1), and then estimating the historical and future climate that the model would give if we could afford to run it at every point within its parameter space. Then we integrate over the parameter space, weighting the projection of future climate at each location according to (a) how likely each combination of parameter values was thought to be before the model simulations were carried out (*prior* information), and (b) the relative likelihood that each point in parameter space gives a true representation of the real climate system (*posterior* information obtained from estimates of how well the model simulates historical climate in practice). This procedure yields probabilities for different outcomes of future climate which are determined by a combination of the complex interactions between physical and biogeochemical processes built into the climate model, expert judgements, structural modelling errors and observational constraints. The interpretation of these probabilities is discussed further in Section 3.3.

Sections 3.2.2–3.2.12 set out a general method for provision of climate projections in any part of the world, at spatial scales skilfully resolved by global climate models (typically regions of approximately 10^6 km² or larger, though this is subject to tests of the validity of its key assumptions as applied in specific regions). However the provision of detailed spatial information for UKCP09 also relies on the addition of a downscaling procedure based on high resolution regional climate model simulations, described in Section 3.2.11. The project was allocated considerable computing resources; however these were inevitably finite, so the methodology relies on judgements regarding how best to deploy these to address the main uncertainties. Assumptions and limitations arising from these choices are highlighted in the following sub-sections.

3.2.2 Process uncertainties

The first task is to define the set of Earth System processes likely to contribute significant uncertainty in 21st century climate (see Box 2.1). These would clearly include surface and atmospheric physical processes (for example water vapour, cloud, surface albedo and soil moisture feedbacks continue to be recognised as key determinants of global and/or regional climate change (Bony *et al.* 2006; Soden and Held, 2006)). However, other components are also likely to be important. Changes in ocean heat transport have potential to influence both global and regional changes (Raper *et al.* 2002; Boer and Yu, 2003), while imperfect knowledge of the radiative forcing due to sulphate aerosols (Anderson *et al.* 2003) is recognised as a significant source of uncertainty, both in determining recent observed climate change and in predicting future changes (Andreae *et al.* 2005). Uncertainties in the fraction of man-made carbon dioxide emissions likely to remain in the atmosphere (due in particular to terrestrial carbon cycle feedbacks) have also emerged as an important source of divergence in future projections by different models, particularly in changes expected during the second half of the 21st century (Cox *et al.* 2000; Friedlingstein *et al.* 2006). We therefore designed our ensemble experiments to sample uncertainties in the atmosphere, ocean, sulphur cycle and terrestrial carbon cycle modules available in the family of HadCM3 components. This covers the major known sources of uncertainty in climate change out to a century or so ahead. Inevitably, however, limitations of computational resource, modelling capability and current understanding imply

that some additional drivers of climate change have to be omitted, or included without sampling of the associated uncertainty. For example, our carbon cycle simulations account for feedbacks associated with ocean as well as terrestrial carbon uptake; however, uncertainties in processes affecting oceanic uptake are not sampled (see Section 3.2.5). Our simulations do not include forcing from carbonaceous aerosols (e.g. Jones *et al.* 2005), non-aerosol atmospheric chemistry (e.g. Johnson *et al.* 2001) or methane cycle feedbacks (Christensen *et al.* 2004; Archer and Buffett, 2005). The sampling of sulphur cycle feedbacks omits the *second indirect effect* arising from the effects of reduced cloud droplet size on precipitation efficiency, and hence cloud persistence, as this process is not included in HadCM3, or indeed in most current climate models (see Table 10.1 of Meehl *et al.* 2007)

Designing ensemble climate projections given finite computing resources

The standard approach to modelling time-dependent climate changes involves simulations which run from pre-industrial conditions up to the end of the period of interest (say from 1860–2100), specifying observed time-dependent changes in external forcing agents (typically man-made changes in greenhouse gases and aerosol precursors, and natural variations arising from solar variability and volcanic eruptions) up to present day, switching to some future scenario of man-made forcings to 2100. The ideal method of sampling modelling uncertainties would be to run a very large ensemble of such *transient* climate change simulations, in which all the relevant Earth System modules (atmosphere, ocean, sulphur and carbon cycle) are coupled together dynamically, and in which different ensemble members sample multiple perturbations to uncertain parameters in all modules simultaneously, in such a way as to ensure comprehensive coverage of the entire parameter space of each module. Such an experiment would ensure that non-linear interactions between all uncertain processes in all modules were thoroughly sampled. Unfortunately, such an experiment is well beyond the available computing resources, so compromises have to be made based on expert judgement of the relative importance of different sources of uncertainty.

Figure 3.1: Elements of our methodology to sample modelling uncertainties using perturbed physics ensembles (PPEs) based on configurations of the HadCM3 climate model. Blue boxes denote ensemble simulations using various model configurations derived from HadCM3. Yellow boxes denote statistical tools required to generate alternative estimates of climate change which combine the sources of uncertainty sampled in the various ensemble experiments. Boxes A and B are described in Section 3.2.3. Boxes C, D and E are explained in Sections 3.2.4, 3.2.5 and 3.2.11 respectively. Boxes F and G represent our timescaling procedure for deriving very large ensembles of realisations of time-dependent climate change from smaller ensembles of climate model simulations, covered in Sections 3.2.4 and 3.2.6. Box H denotes our downscaling procedure (see Section 3.2.11) for the generation of probabilistic projections at the 25 km resolution required for UKCP09, derived from information at larger scales obtained from global climate model simulations.

Sampling uncertainties with perturbed physics ensembles

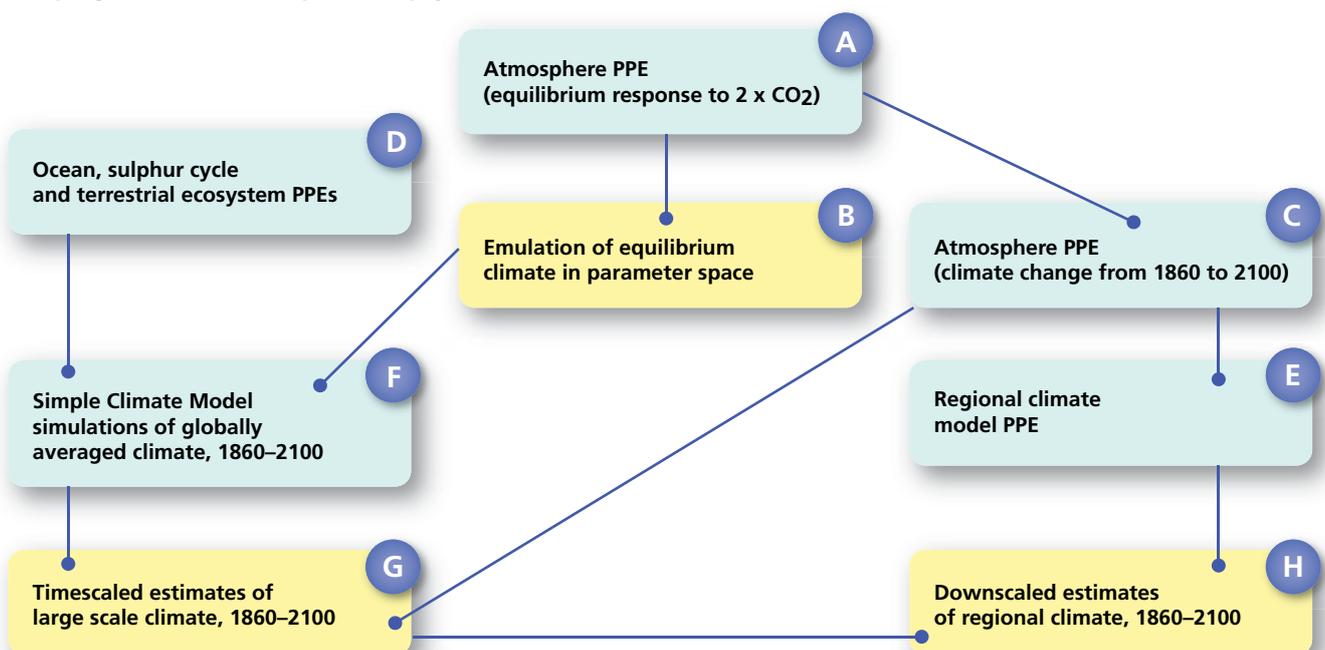


Figure 3.1 gives a schematic summary of the major components of our strategy for sampling modelling uncertainties, through the combination of a number of ensemble climate projection experiments. These experiments use several model configurations derived from HadCM3 to sample uncertainties in climate change during the 21st century, and are described below in Sections 3.2.3–3.2.6, and 3.2.11.

3.2.3 Sampling uncertainties in surface and atmospheric processes

Based on the assessment that surface and atmospheric feedbacks are likely to provide the largest source of uncertainty in regional changes during the coming century, we focus our resources on sampling the parameter space of these processes more comprehensively than those of the ocean, sulphur cycle or carbon cycle modules. The atmosphere module of HadCM3, which also includes land surface processes and surface–atmosphere exchanges, contains 100 or more parameters controlling the model parameterisations of small scale processes (which cannot be resolved explicitly on the model grid) in terms of grid box variables. It would not be computationally feasible to explore the combined effects of perturbing all these parameters, and in any case some parameters exert a much more significant influence than others on the simulated outputs of the model. Parameterisation experts were therefore asked to identify a subset of these which control the main processes most important for the simulation of (both global and regional) climate, and then to estimate plausible minimum, intermediate and maximum values (accepting that, in general, there would be insufficient evidence to provide a unique specification of the likely distribution of parameter values between the minimum and maximum values). This exercise resulted in a subset of 31 key parameters for perturbation. We assume that neglect of possible perturbations to additional parameters does not significantly affect the spread of model behaviour generated from our simulations.

Simulations of equilibrium climate changes in response to doubled CO₂

A large ensemble of (at minimum) a few hundred members is required to provide a reasonable first-order estimate of how the model behaviour varies within this 31-dimensional space, given that both the linear effects of each parameter (Murphy *et al.* 2004), and non-linear interactions between them (Stainforth *et al.* 2005), can have important influences on the model simulations. Resource limitations prevented us from undertaking ensembles of transient climate change simulations of this size, so the required large ensemble was run using a computationally less demanding model configuration (HadSM3) in which the atmosphere module is coupled to a simple thermodynamic model of the near-surface ocean, which warms or cools in response to surface heat exchanges with the atmosphere, and in which horizontal and vertical transport within the ocean is prescribed. Such a model configuration is widely accepted as a suitable set-up for the simulation of equilibrium climate changes, including the climate sensitivity, a standard benchmark of climate change defined as the global mean equilibrium response of surface temperature to doubled carbon dioxide. However, this simplified approach neglects climate change feedbacks involving changes in regional ocean heat transport (Boer and Yu, 2003), and implies the need for a method of converting simulated equilibrium changes into corresponding estimates of transient climate change. This conversion relies on the assumption that a reasonable relationship exists between patterns of time-dependent and equilibrium climate changes in response to increasing greenhouse gas concentrations. Harris *et al.* (2006) find a close relationship for multiyear averages of surface temperature changes, whereas for precipitation the degree of correspondence varies significantly with location, though it is quite good for the UK and Europe. Note, however, that our conversion method (described in

Section 3.2.4) also accounts for random and systematic differences between simulated patterns of time-dependent and equilibrium changes.

An ensemble of 280 HadSM3 experiments was run, sampling the effects of perturbing these parameters relative to the settings used in the standard published variant of HadCM3 (Gordon *et al.* 2000). These settings are referred to hereafter as the *standard* parameter values, though a number of these values actually correspond to extremes of the ranges identified by experts, due to the practice of *tuning* the model to improve its simulation of certain basic aspects of climate, such as the planetary radiation balance. Each experiment consisted of a *control* simulation of recent climate, and a simulation of the response to a doubled carbon dioxide concentration, run for a sufficient length of time to allow the resulting climate change to reach equilibrium. Murphy *et al.* (2004) carried out an initial ensemble of 53 members in which one parameter was perturbed at a time. This was subsequently augmented by a second ensemble of 128 members containing multiple parameter perturbations chosen to sample a wide range of climate sensitivities, achieve skilful simulations of present climate and maximise coverage of parameter space (details in Webb *et al.* 2006). Further HadSM3 simulations were then run to achieve improved sampling of parts of parameter space influenced by key interactions between parameters (Rougier *et al.* 2008). Together, these ensembles provide the 280 simulations used in UKCP09.

Emulation of equilibrium climate changes in response to doubled CO₂

This set of simulations is sufficient to sample the main effects of parameter variations within our 31-dimensional space, but not to cover it comprehensively. We therefore use a statistical tool called an emulator (e.g. Rougier *et al.* 2008), to help us estimate the values of the required set of climate variables at any given point in parameter space. The emulator is trained on the available GCM simulations to estimate the results of a set of historical and future climate variables required in the production of our probabilistic projections. Each climate variable is emulated using an equation which provides a best estimate value and associated errors for any combination of model parameter values. This is done by using the available GCM simulations to train multiple regression relationships which express the required climate variables as functions of the model parameters, where the set of regressors capture key interactions between the effects of different parameters, as well as the effects of each parameter in isolation. Emulation errors are guaranteed to be greater than or equal to internal climate variability, and are typically 20–50% larger.

Using the emulator, we are then in a position to integrate over the whole of our parameter space, estimating values of both historical climate variables (required to weight each location according to how well the GCM would simulate historical climate given that particular combination of parameter settings), and future climate changes. This integration allows us to estimate observationally constrained probabilities for different changes, accounting for model uncertainties. It provides the bedrock of our approach to probabilistic projection; however, a number of additional elements are required to convert the results into user-relevant estimates of climate change for specific 21st century periods, and to ensure that additional sources of uncertainty are included. These are described in Sections 3.2.4–3.2.11. Several aspects of the methodology (in addition to the emulation stage described here) require the estimation of uncertainties from the residual errors of statistical regression or optimisation procedures. These statistical errors are assumed to be Gaussian, and they are all included in the uncertainty expressed in the projections. In view of this, several of the UKCP variables are transformed prior to the calculation of projected changes, the

inverse transformation being applied afterwards to recover projected changes in the original variables. These transformations are made either to reduce the risk of non-Gaussian error characteristics, or to ensure that absolute bounds in some of the projection variables cannot be exceeded by the addition of several sources of statistical error. In particular, this ensures that variables presented as percentage changes relative to the UKCP baseline period cannot go beyond -100%.

3.2.4 Sampling uncertainties in transient climate change

The experiments described in Section 3.2.3 provide estimates of the equilibrium climate change in response to doubled carbon dioxide, which must be converted into estimates of 21st century changes. This is done by running a smaller ensemble of simulations of transient climate change, in which the atmosphere module is coupled to the full three-dimensional ocean module of HadCM3, which simulates horizontal and vertical transport processes dynamically. The configuration of HadCM3 for these experiments is as described by Gordon *et al.* (2000), except that the representation of the atmospheric sulphur cycle is upgraded to use the fully interactive module of Jones *et al.* (2001), thus avoiding the need to approximate the effect of sulphate aerosol on cloud albedo using an offline calculation (Johns *et al.* 2003).

The approach involves a 17 member ensemble (PPE_A1B) which samples a subset of the atmospheric module parameter combinations used in the larger HadSM3 ensemble described above. One member used the standard HadCM3 parameter settings, the sixteen additional members using combinations of perturbed settings chosen to sample a wide range of climate sensitivities, while also sampling a wide range of alternative parameter values and providing credible simulations of historical climate. Flux adjustments are used to limit simulation biases in sea

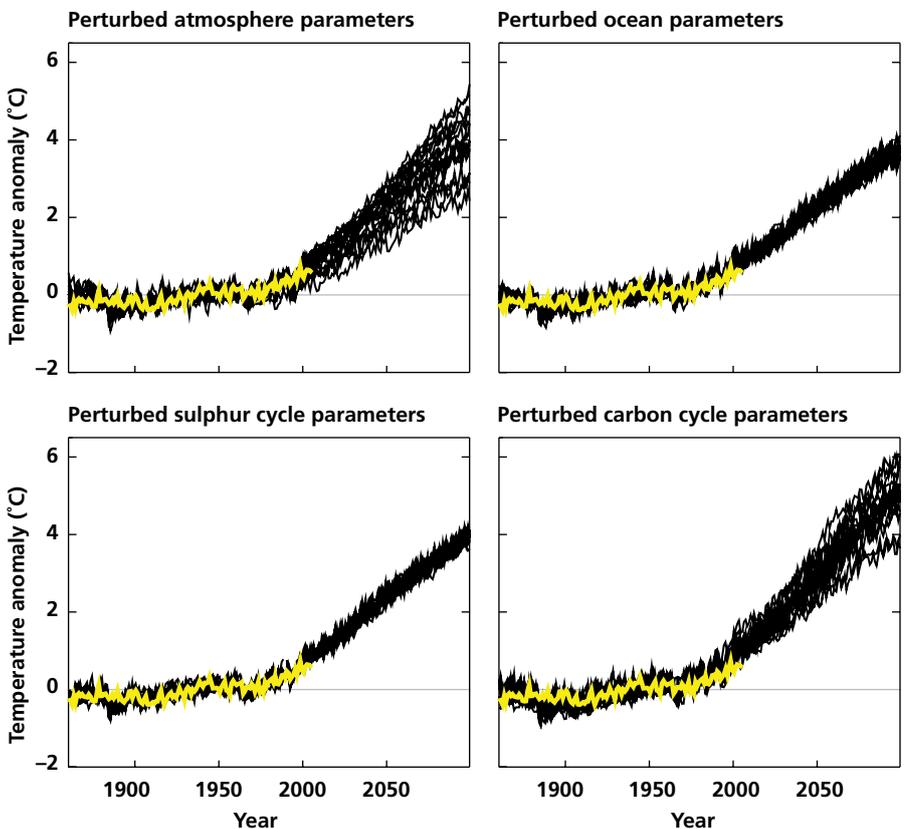


Figure 3.2: Global, annual mean 1.5 m temperature anomalies (°C) from different perturbed physics ensembles of time-dependent climate change under SRES A1B emissions, from 1860 to 2100. Anomalies are expressed with respect to the 1860–2000 mean. Each plot shows observations in yellow, with ensemble projections in black. Top left: Ensemble PPE_A1B, sampling perturbations to atmosphere model parameters. Top right: Ensemble with perturbations to ocean model parameters. Bottom left: Ensemble with perturbations to sulphur cycle parameters. Bottom right: Ensemble with perturbations to terrestrial ecosystem (carbon cycle) parameters.

surface temperature and salinity. The sampling of parameter space and climate sensitivity, and the calculation of flux adjustments, was based on (but updated from) an earlier PPE of HadCM3 variants described by Collins *et al.* (2006). Perturbed model variants in PPE_A1B give global simulations of historical climate of comparable quality to the standard model variant, as was also found in the Collins *et al.* (2006) experiment; however, improvements to the flux adjustment technique in PPE_A1B removed biases in sea surface temperature and salinity found in the North Atlantic and Arctic Oceans in the simulations of Collins *et al.* By reducing regional systematic errors the flux adjustment process helps to ensure that the ensemble projects credible regional climate changes, and it also allows the effects of parameter perturbations on the transient response to be explored without being excessively constrained by the need to achieve precise balance in the planetary radiation budget. The simulations were started in the year 1860, and driven up to 2000 by historical time series of concentrations of greenhouse gases (carbon dioxide, methane, nitrous oxide, chlorofluorocarbons and ozone), sulphur emissions, and reconstructions of variations in solar activity and volcanic aerosol. From 2000 to 2100 they were driven by future concentrations of greenhouse gases and sulphur emissions from the SRES A1B scenario. The results show a substantial spread in projections of future global temperature rise (Figure 3.2). Here, and in Sections 3.2.5–3.2.12 we describe the entire methodology as applied in the case of the A1B scenario. Extensions to cover the A1FI and B1 scenarios are summarised in Section 3.2.13.

Estimating transient changes from equilibrium changes using timescaling

While these 17 transient simulations provide a limited sample of direct realisations of time-dependent climate change, our methodology requires that we estimate the time-dependent response from any point in the model parameter space referred to above. This is achieved by developing relationships between the equilibrium response of HadSM3, and the transient response of HadCM3, using the PPE_A1B HadCM3 simulations and the 17 member subset of the larger HadSM3 ensemble containing corresponding parameter perturbations to the PPE_A1B members. Once calibrated, these relationships can then be used to estimate the regional transient response of relevant climate variables (see Table 1.1) that would be obtained with any desired combination of parameter settings, thus providing the basis for the generation of probabilities for regional, time-dependent climate change through the integration over model parameter space referred to above.

The method, which we term *timescaling*, has been developed from earlier work by Harris *et al.* (2006): It involves normalising the regional equilibrium response of HadSM3 simulations by their climate sensitivities, and then scaling the normalised response according to the transient response of global average surface temperature, which is simulated using a simple climate model tuned to the climate sensitivity of the relevant ensemble member. The simple model is based on that of Rowntree (1998) and simulates globally-averaged land and ocean surface temperatures in response to imposed radiative forcing anomalies, representing vertical heat transfer in the ocean via a globally averaged heat diffusion equation, modified to include upwelling and downwelling following Schlesinger *et al.* (1997). This procedure provides time-dependent estimates of regional climate change, which are modified by a correction term (also scaled according to global mean temperature) which allows for differences between the characteristic patterns of equilibrium and transient climate change arising from the effects of oceanic thermal inertia and changes in ocean circulation. In principle the correction term is liable to depend on the values of the model parameters; however, we neglect such dependencies as we do not possess enough

transient HadCM3 simulations to quantify them robustly. Also, this approach will not be able to replicate time-dependent responses which are non-linearly related to changes in global mean temperature, for example over northern Australia, where precipitation initially increases with global temperature in our perturbed physics simulations, but later reduces as the global response becomes larger (Harris *et al.* 2006). Over the UK, we do not see evidence of substantial non-linearities of this nature. However the method does include an error term which captures bias and uncertainty in our timescaled estimates of regional changes. This adjusts the projections to allow for the estimated effects of errors associated with our assumption in that the transient response is linearly related to global temperature, and also accounts for the effects of internal climate variability, errors in our simple climate model projections of the global temperature response found in HadCM3 simulations, and our assumption that the correction term is invariant across parameter space. It is assumed to take the form of a Gaussian distribution, noting that some variables are transformed to ensure that this assumption is reasonable (see Section 3.2.3). The time-dependent means and variances of these distributions are calculated by using the PPE_A1B simulations to verify timescaled estimates derived from equilibrium changes simulated by HadSM3 ensemble members containing corresponding sets of parameter perturbations. The correction term is also obtained in this fashion.

The timescaling process is illustrated by Figure 3.3(a), which shows projections of summer temperature changes over the global climate model grid box corresponding to Wales from the 17 HadCM3 projections (left panel), compared against corresponding timescaled projections (right panel). The coloured lines in the right panel represent projections obtained by scaling the relevant HadSM3 equilibrium responses according to global mean temperature, and adding the correction term accounting for differences between the characteristic patterns of equilibrium and transient climate change (see above). These lines can be interpreted as estimates of the forced transient component of climate change, in the absence of non-linear dependencies of the forced response on global mean

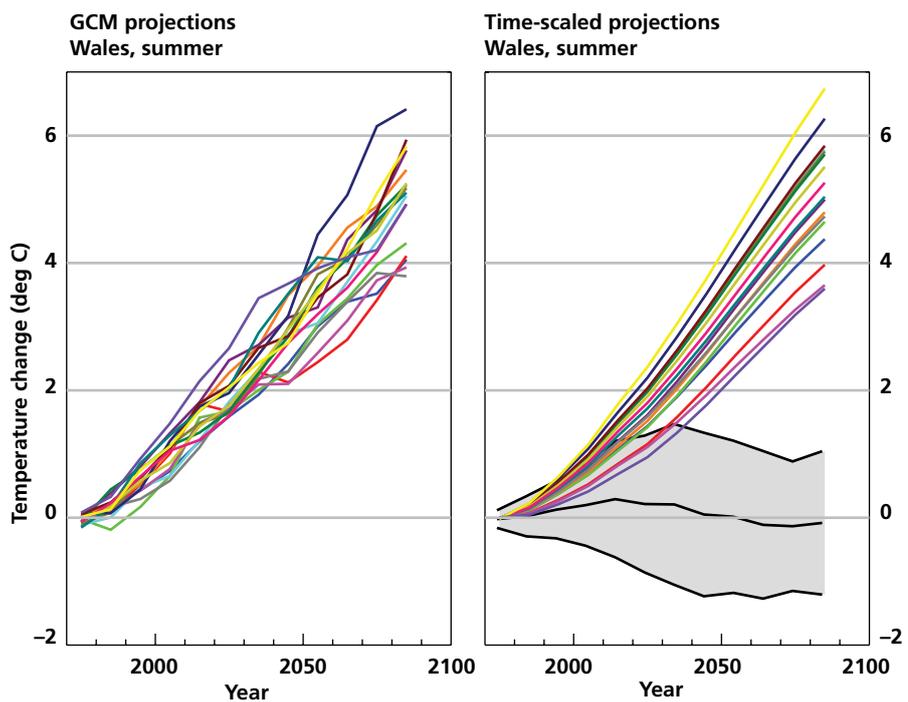


Figure 3.3(a): Left panel shows projected changes in 30-yr averages of surface temperature (°C) relative to 1961–1990 over the global climate model grid box corresponding to Wales, in summer, for the 17 members of the PPE_A1B ensemble of perturbed HadCM3 variants. Right panel shows estimates of the changes derived from the timescaling procedure described in the text (coloured lines). The grey shading illustrates the range of timescaling uncertainties, defined as plus and minus two standard deviations of the errors found by timescaling each of the 17 HadCM3 projections in turn, using statistics obtained by calibrating the procedure using equilibrium and time-dependent climate changes from the other 16 ensemble members.

temperature. In this case, the envelope of timescaled projections corresponds quite closely with that defined by the climate model projections. However the smoothed coloured lines of the timescaled estimates deviate in detail from their climate model counterparts at any given time period, due to the effects of internal variability, non-linear dependencies on global temperature, and other uncertainties in the timescaling process. For this reason, the order of the coloured lines in the timescaled estimates differs somewhat from their HadCM3 counterparts, at any given time level. However, the effects of these timescaling errors (shown separately as grey shading in Figure 3.3(a)) are included in the UKCP09 projections as described above, by adding time-dependent uncertainties sampled from our error estimates to the basic timescaled projections shown by the coloured lines. Results for winter precipitation changes (Figure 3.3(b)) are similar in character, except that the envelope of climate model projections is significantly wider than that of the timescaled projections out to about the 2050s. This is mainly because the forced climate change for the next few decades (estimated in isolation by the coloured lines in the right panel) is relatively small compared to the component of the spread in the climate model projections explained by internal variability. However, we emphasise that the timescaling error term (grey shading) does capture the effects of internal variability, so this component of uncertainty is included in the full envelope of timescaled projections (not shown in Figures 3.3(a) and (b), but obtained by combining the coloured lines and the grey shading).

While changes in well-mixed greenhouse gases such as carbon dioxide give rise to spatially uniform changes in radiative forcing, this is not the case for other forcing agents included in our transient simulations (historical and future changes in sulphate aerosols and ozone, and historical changes in solar and volcanic activity). The forcing due to sulphate aerosols, in particular, is concentrated over and downstream of industrialised regions of the northern hemisphere (Forster *et al.* 2007). The patterns of climate change in response to spatially heterogeneous forcings cannot be assumed to follow that found in response to well-mixed greenhouse gases. We account for this by running an additional

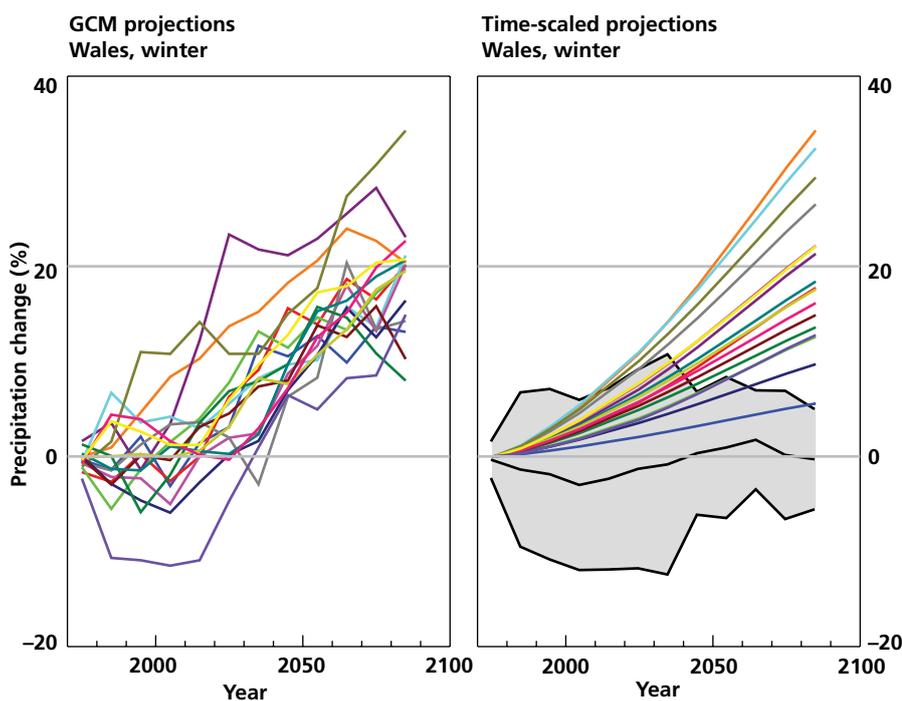


Figure 3.3(b): As Figure 3.3(a), but for precipitation changes (%) over Wales in winter.

17 member perturbed physics ensemble of HadCM3 simulations from 1860 to 2100, identical to PPE_A1B except that concentrations of well-mixed greenhouse gases are held fixed at pre-industrial levels, allowing the climate response to the heterogeneous aspects of the forcing to be isolated. Results from this ensemble (PPE_A1B_NOGHG) are used to estimate the regional response to these forcings (per unit global temperature change) as a function of time, which then forms a potential additional contribution to our timescaled estimates of transient climate change. We do not possess sufficient simulations to estimate how the normalised response to heterogeneous forcing agents might vary across the model parameter space. However, future changes in forcing in the emissions scenarios considered for UKCP09 are dominated by well-mixed greenhouse gases, and for these we do estimate variations across parameter space in greater detail.

In practice, the added refinement of including a separate term for the heterogeneous forcings is found to be important for some variables, but not others. Use of this term is therefore determined on a case-by-case basis, dependent on whether its inclusion leads to a statistically significant reduction in the uncertainty associated with our climate change estimates.

3.2.5 Sampling uncertainties in additional Earth System processes

We sample uncertainties in ocean, sulphur cycle and terrestrial carbon cycle processes by running three additional perturbed physics ensembles, each consisting of 16 perturbed variants of HadCM3. Each of these ensembles is driven from 1860 to 2100 by the same time series of forcing agents used in PPE_A1B. In each of these ensembles parameters in the module targeted for perturbation are varied within ranges obtained by consultation with experts, while parameters in other modules are held fixed at values used in the standard model variant. In all cases parameter combinations were determined using a Latin Hypercube sampling design (McKay *et al.* 1979).

Ocean transport

The first ensemble addresses uncertainties in ocean transport, building on preliminary simulations reported by Collins *et al.* (2007). The ensemble members sample perturbations to parameters controlling various aspects of the resolved and subgrid-scale transports of heat, salt and momentum in both the horizontal and vertical. In these simulations, future global mean temperature rise shows a limited dependence on these ocean parameters (Figure 3.2), much smaller than the uncertainties arising from atmospheric processes.

Sulphur cycle

The second ensemble samples uncertainties in atmospheric sulphur cycle processes, represented in HadCM3 using the module described by Jones *et al.* (2001). It simulates sulphate aerosol concentrations from prescribed emission fields of anthropogenic sulphur dioxide (SO₂), natural dimethyl sulphide and tropospheric sulphur arising from quasi-regular volcanic eruptions. Three modes of aerosol are represented, comprising sulphate dissolved in cloud droplets plus two free particle modes. The model simulates production of sulphate by oxidation of SO₂, transport within the atmosphere, rain out and transfers between the different aerosol modes. The atmospheric sulphur burden affects radiation via the direct (cooling) influence of scattering and absorption of incoming solar radiation, and through increases in cloud albedo resulting from the action of sulphate aerosols as cloud condensation nuclei (the *first indirect effect*). As mentioned earlier, the *second indirect effect*, in which reductions in cloud droplet size reduce precipitation efficiency and increase cloud lifetime, is not included since the calculation of precipitation in HadCM3 does not allow for any

dependence on cloud droplet number concentrations. The 16 member ensemble of HadCM3 simulations samples simultaneous perturbations to parameters controlling key aspects of the processes outlined above, including emissions of aerosol precursors. All ensemble members used the settings for atmosphere and ocean module parameters employed in the standard variant of HadCM3. This ensemble simulates a wide range of atmospheric sulphur burdens (although perturbations to some of the atmosphere module parameters in PPE_A1B and PPE_A1B_NOGHG also have a significant impact on these). The impact of sulphur cycle perturbations on global mean temperature changes is modest compared with that in PPE_A1B (Figure 3.2).

Terrestrial ecosystem

Uncertainties in terrestrial ecosystem processes are sampled in a third ensemble in which the TRIFFID dynamic vegetation module of Cox (2001) is added to HadCM3, to form an Earth System model HadCM3C. TRIFFID simulates soil carbon, and the growth and replacement of five functional types of vegetation (broadleaf tree, needleleaf tree, C3 grass, C4 grass and shrubs). The functional types vary according to the net available carbon and competition between plant types, parameterised using empirical relationships. Soil carbon can be increased by litterfall and is returned to the atmosphere by microbial respiration, which depends on temperature and soil moisture. CO₂ fluxes at the land–atmosphere interface are determined by photosynthesis and plant and microbial respiration. In order to simulate carbon fluxes at the ocean–atmosphere interface, an ocean carbon cycle module (Cox *et al.* 2001) is also included. This simulates exchange of gaseous CO₂ with the atmosphere, the transport of dissolved inorganic carbon and cycling of carbon by marine biota via a nutrient–phytoplankton–zooplankton–detritus ecosystem module (Palmer and Totterdell, 2001) that accounts for the effects of light penetration, alkalinity and nutrient availability on biological carbon uptake. In previous carbon cycle experiments using HadCM3 (e.g. Cox *et al.* 2000; Jones *et al.* 2003), the horizontal resolution of the ocean module was reduced; however, here we maintain the standard resolution of 1.25 x 1.25 degrees in order to ensure that our carbon cycle simulations are physically consistent with the other coupled ocean–atmosphere ensembles included in our methodology.

A 16-member ensemble was produced, sampling simultaneous perturbations to TRIFFID parameters controlling soil carbon and the five vegetation functional types. A further ensemble member with standard TRIFFID settings was also run. Parameters in the ocean carbon cycle module were held fixed at standard values in these simulations, because resource and time limitations made it impractical to perform the ensemble of long preliminary integrations (e.g. Cox *et al.* 2001) which would have been required to achieve equilibrium in ocean–atmosphere carbon fluxes. The specification of forcing agents was as in PPE_A1B, except that CO₂ was input as a time series of emissions rather than concentrations, in order to allow carbon cycle feedbacks to operate. This ensemble simulates a substantial range of future changes in CO₂ concentration (669–1130 ppm at the year 2100, for example), and therefore of global mean surface temperature (Figure 3.2), comparable to the spread found by sampling physical surface and atmospheric processes in PPE_A1B. Uncertainties in the ocean carbon sink are not sampled in our simulations (as explained above); however, the spread of responses obtained is similar to that found in a previous multi-model ensemble of carbon cycle simulations carried out in the Coupled Climate Carbon Cycle Intercomparison Project (C⁴MIP) by Friedlingstein *et al.* (2006). The C⁴MIP ensemble sampled variations in both terrestrial and ocean carbon cycle processes and found that climate-induced changes in carbon storage were explained mainly by the former.

In addition to their impacts on global mean surface temperature (Figure 3.2), the ocean, sulphur cycle and terrestrial ecosystem PPEs all show some statistically significant impacts on patterns of regional change in some parts of the world. For example, the sulphur cycle PPE shows a significant spread in temperature changes in the Arctic Ocean and over interior regions of the northern Eurasian landmass (because surface albedo feedbacks amplify the effects of perturbations to the response of surface temperature), and in precipitation changes over tropical regions of the central and western Pacific Ocean (due to the strong coupling with sea surface temperature changes in these regions). The ocean PPE shows similar impacts over the Arctic and tropical Pacific Oceans, while the terrestrial carbon cycle PPE shows a large spread of precipitation changes over Amazonia, due to the regional influence of ecosystem-atmosphere interactions (Betts *et al.* 2004). However the impacts on changes over the UK (beyond those directly explained by variations in the global mean warming) turn out to be relatively minor.

3.2.6 Combining uncertainties in different Earth System processes

The Earth System ensembles described in Section 3.2.5 are not large enough to provide a basis for training an emulator capable of estimating the model response at any point in the parameter space of ocean, sulphur cycle or carbon cycle processes (cf Section 3.2.3). This prevents us from including the relevant uncertainties via a formal application of Bayes theorem in an integration over the model parameter space (cf. Section 3.2.7 below). However, we do include uncertainty estimates obtained from these ensembles in a simpler manner, by generalising the timescaling technique described in Section 3.2.4. This is done by configuring the simple climate model used in timescaling to include sulphate aerosol forcing, and simple globally averaged parameterisations of processes associated with the effects of terrestrial carbon cycle feedbacks on the atmospheric CO₂ concentration. When running the simple model to estimate the transient climate response for some specified set of surface and atmospheric HadCM3 parameters, we sample the effects of additional Earth System processes by selecting from a distribution of possible values for the simple model parameters controlling global mean ocean heat uptake, sulphate forcing or CO₂ concentration. For heat uptake, this is done by calculating values of ocean diffusivity for each of the 17 members of our ocean perturbed physics ensemble (Section 3.2.5), and also from 20 alternative simulations derived from the multi-model ensemble of coupled ocean-atmosphere models submitted to the IPCC AR4. The multi-model ensemble values were taken from the 23 models listed in Table 8.1 of Randall *et al.* (2007), omitting two models because data required for the calculation were not available, and one because the wrong climate change forcing was applied in the relevant experiment. Inclusion of the multi-model ensemble results enabled us to account in a simple way for structural uncertainties in ocean transport processes not sampled in our perturbed ocean ensemble. Values are then selected from these 37 possible values, assuming each to be equally plausible.

Including sulphate aerosol forcing uncertainties in timescaled projections

For sulphate aerosol forcing the approach is somewhat more complicated, because variations in physical atmospheric parameters (particularly those associated with cloud) are found to exert a significant influence on the forcing, in addition to variations in parameters directly associated with the sulphur cycle. Furthermore, a significant relationship between global mean aerosol forcing and climate sensitivity was found in our PPE_A1B_NOGHG ensemble (low sensitivity model variants tend to simulate high levels of low cloud, and therefore simulate larger changes in forcing in response to aerosols). We accounted for these factors by developing a regression relationship between a transformed function of aerosol forcing, and global climate feedback (the reciprocal of climate sensitivity). The

distribution of forcing values is Gaussian in the transformed units. Variations in transformed aerosol forcing, diagnosed from the 16-member perturbed sulphur cycle ensemble, were assumed independent of atmospheric perturbations and added to each member of our PPE_A1B_NOGHG ensemble, thus forming a dataset for regression which sampled uncertainty arising from both atmospheric and sulphur cycle processes. When running the simple climate model for a given location in parameter space (and hence a given climate sensitivity), we then sampled alternative aerosol forcing values from the error statistics of the regression relationship. This method gives a distribution of aerosol forcing values for present day climate (relative to pre-industrial conditions) similar to that given in the IPCC AR4 (see Figure 2.20 of Forster *et al.* 2007), based on the statistical assessment of the uncertainty of radiative forcing mechanisms documented by Haywood and Schulz (2007).

Including carbon cycle feedback uncertainties in timescaled projections

Given that carbon cycle uncertainties provide a leading order contribution to the uncertainty in global mean changes, and recognising that our perturbed physics ensemble does not sample uncertainties associated with structural carbon cycle assumptions in HadCM3C, we also include results from the C⁴MIP multi-model simulations in our sampling of possible feedbacks. We performed a pre-screening exercise in which the historical simulations of global carbon budget components (fraction of anthropogenic emissions stored in atmosphere, land and ocean) were compared with an observational constraint based on records of atmospheric CO₂ increase, estimates of total emissions (fossil fuel plus land use emissions) and the oceanic uptake of anthropogenic CO₂ (Sabine *et al.* 2004). Two of the perturbed physics simulations and one of the C⁴MIP simulations were found to be inconsistent with the spread of plausible values implied by estimates of observational uncertainty, so these were excluded. We also excluded results of the HadCM3LC model contributed to C⁴MIP, as this model is strongly related to that used for our perturbed physics simulations. This left 9 members of the C⁴MIP ensemble and 15 members of the perturbed physics ensembles, whose simulated global mean feedbacks were sampled in the timescaling procedure, assuming all 24 estimates to be equally plausible.

The parameterisation of carbon cycle feedbacks in the simple climate model contains explicit temperature dependences, allowing the (significant) effect of variations in the global temperature response on the global mean carbon cycle response to be captured (e.g. Andreae *et al.* 2005). This is achieved using globally averaged calculations of changes to the vegetation and soil carbon stores consistent with the main features of the corresponding calculations used in the terrestrial ecosystem module of HadCM3 (Jones *et al.* 2003), which contains temperature-dependent parameterisations of photosynthesis and plant and soil respiration. With the exception of this carbon cycle–temperature relationship, and the aerosol forcing–climate sensitivity relationship described above, our timescaling method does not account for non-linear interactions between the global feedbacks in different Earth System modules. This is because time and resource limitations prevented us from running HadCM3 ensemble simulations in which parameters in all component modules were varied simultaneously. The UKCP09 projections are conditional on the assumption that additional non-linear interactions are likely to be small compared with the two significant known relationships referred to above. This issue is a subject of current research.

Potential contributions of ocean, sulphur cycle and carbon cycle processes to uncertainties in regional climate changes (beyond the effects directly attributable to uncertainties in global mean surface temperature) are not accounted for in

the generalised timescaling technique. This is because results from the relevant ensembles indicate that such contributions would be relatively minor for changes over the UK (Section 3.2.5), and also because quantification of the impacts of non-linear interactions is beyond the scope of the experimental design for UKCP09 (see above). In some regions neglect of such regional effects would not be realistic, a good example being Amazonia where carbon release from forest dieback is dependent on regional changes in precipitation (Betts *et al.* 2004). The extent to which the UKCP methodology could be applied in other parts of the world will therefore depend upon careful evaluation of the potential impacts of regional effects not covered by our timescaling procedure, in addition to the validity of further assumptions required by our technique, such as the use of a linear scaling to global mean temperature changes (see Section 3.2.4).

3.2.7 Probabilistic projections of the equilibrium response to doubled CO₂

In Sections 3.2.7–3.2.9 we describe how we obtain probabilistic projections for the equilibrium response to doubled CO₂ concentration. This exercise provides marginal probabilities for changes in individual variables, or joint probabilities for changes in two or more variables (e.g. temperature and precipitation in some specific region), at the spatial scale of HadSM3 grid boxes (approximately 300 x 300 km²). However it is also necessary to apply our timescaling procedure (Sections 3.2.4 and 3.2.6), and the downscaling procedure (described in Section 3.2.11 below), to obtain estimates of 21st century changes at the local scales required by UKCP09 users. The combination of these elements is outlined later, in Section 3.2.12.

Probabilistic projections are obtained using the Bayesian statistical framework introduced in Section 3.2.1, described here in general terms, omitting technical details. The calculation is based on values of variables of historical and future climate obtained from a climate model whose outputs depend upon a set of parameters controlling processes judged to be important determinants of the quality of its simulations. Observed values of the historical variables and their associated errors are also required, in order to weight model outputs according to their quality. Probabilities for different values of future variables are obtained by applying Bayes Theorem through an integration over the model parameter space of surface and atmospheric processes (henceforth referred to as *parameter space*), as outlined in Section 3.2.1 (see Rougier (2007) for more details). However, we cannot afford to run the climate model itself at every point within this space, so we train an emulator to replicate the model outputs (see Section 3.2.3), and then use the emulator to estimate values of the required variables for any given combination of parameter settings.

The Bayesian framework allows (and requires) us to account for relationships between the various elements involved in the calculation. Some simplifying assumptions are necessary to make the calculation tractable: for example there is no obvious reason to expect that errors in emulated estimates of climate model output would be correlated with errors in observed estimates of the true historical climate, so we assume these to be independent. On the other hand, our method relies on the basic assumption that relationships can be found between variations across parameter space in the modelled values of historical climate and future changes (e.g. Piani *et al.* 2005; Knutti *et al.* 2006), so we would want to account for these in the calculation. In our Bayesian approach, this is achieved by calculating weights for different combinations of parameter values according to how well the model simulates a set of historical observations given those values. These posterior weights constrain the model parameter space to regions

giving rise to relatively skilful simulations, and thus also constrain projections of future climate variables, to an extent which depends on how strongly the future variables are controlled by values of model parameters. This helps to reduce the dependence of the projections on expert prior choices imposed by the experimenters (see Annex 2). Also, the simulated changes, and their associated uncertainties, can be adjusted according to the errors in the simulated values of historical observables, according to the strengths of the correlations between them. This ability to pick out key relationships from a range of possible influences is a critical strength of the procedure, because future changes in climate over the UK (indeed in any region) are influenced by an array of feedback processes, some of which are local in origin, and some of which involve remote influences. Rowell and Jones (2006) demonstrate this in relation to future summer drying over Europe, for example, showing that this is affected by large scale thermodynamic feedbacks, changes in atmospheric circulation, and regional changes in soil moisture influenced by surface–atmosphere coupling in summer, and also by changes in the annual cycle of surface hydrological components dependent on changes in temperature, snowmelt and precipitation at other times of the year. Thus it would not be possible to determine the credibility of projected future changes by focusing solely, for example, on simulated values of historical metrics limited to the region and season of interest (e.g. Moberg and Jones, 2004). The set of observations used to constrain the UKCP09 projections is described in Section 3.2.9.

The complex and interconnected nature of changes in different variables (illustrated by the example above) also suggests that it would be difficult to justify assigning different weights to projections of different variables from the same model variant. Our statistical framework reflects this, being based on the assumption that each model variant should be assigned a universal weight which reflects the quality of its ability to simulate climate as a whole. This weight quantifies the relative likelihood that a given combination of parameter settings provides a representation of climate system processes consistent with our observations of the real world. The likelihood depends on the difference between the emulated values of our set of historical variables and the corresponding observations, accounting for covariances between the variables, and normalized by the uncertainty in the differences, obtained by adding contributions from emulator error, observational error and structural modelling error. The sizes of the covariances determine how rapidly the weight drops as the emulated values move away from observations. The structural error arises from the recognition that HadCM3 (like any climate model) contains certain fundamental biases which cannot be resolved by varying its uncertain parameters, so the framework includes a key term called *discrepancy* which captures the additional uncertainty in model projections arising from such errors.

In our integration over model parameter space, we assume that climate model parameters are *a priori* equally likely within the middle 75% of the range estimated by experts, and that the probability drops linearly to zero at the minimum and maximum values. It is recognised that alternative and equally defensible prior distributions could be proposed (e.g. Rougier and Sexton, 2007); however, the results are quite robust to a number of reasonable alternative choices (see Annex 2).

3.2.8 Structural model errors (*discrepancy*)

What is discrepancy, and why is it important?

The discrepancy term, introduced in Section 3.2.7, is a measure of how informative the climate model is about the real world. Formally, it represents the mismatch we would find between the model and the real world if we could locate precisely the

combination of model parameter settings giving the best overall simulation of climate that the model is capable of providing. Discrepancy applies to simulations of both historical and future climate. It is also a prior input to the Bayesian framework, and should therefore be specified using a method as independent as possible from the specific observations used to weight the (emulated) climate model projections, in order to avoid double counting the observations. Discrepancy is itself uncertain, and is therefore specified as a distribution (in common with other uncertain inputs to the Bayesian calculation). Values must be specified for all historical and future variables involved in the calculation, including covariances between the variables. Discrepancy in historical variables focuses the weight on the well modelled variables and prevents small variations in the poorly modelled variables from having an unduly large impact on the weighting. Discrepancy in future variables increases the uncertainty associated with the projections, and mitigates the risk of making overconfident projections. Specifying the discrepancy is an extremely demanding task in principle, given the inherent difficulty of anticipating the effects on particular climate variables of missing or inadequately understood processes, and their complex interactions.

Estimation of discrepancy in UKCP09

In practice we estimate discrepancy by using results from our large ensemble of HadSM3 simulations of present day and doubled CO₂ climates (see Section 3.2.3) to predict the results of an ensemble of different climate models, whose members consist of coupled atmosphere–mixed layer ocean (*slab*) models of similar complexity and credibility as HadSM3, but employing different basic assumptions in some of their parameterisations of physical processes. Note that this exercise must be carried out using ensembles of slab model simulations, rather than ensembles of coupled models containing a full dynamical ocean (e.g. Figure 3.2), because our perturbed physics ensembles using HadCM3 are too small to support a direct application of the Bayesian framework to their results. Nevertheless, our approach confers the benefit of allowing us to provide projections which combine results from perturbed physics and multi-model ensembles, hence adjusting the projections to account for likely biases arising from structural errors in HadCM3. It is based on the judgement that the effects of structural differences between models can be assumed to provide reasonable *a priori* estimates of possible structural differences between HadSM3 and the real world. We take a given multi-model ensemble member as a proxy for the true climate, and use our emulator of HadSM3 to locate a point in the HadSM3 parameter space which achieves the best multivariate fit between HadSM3 and the multi-model member, based on a set of climate variables described in Section 3.2.9. The fit is determined using an optimisation procedure starting from a randomly-selected initial point in parameter space. The difference represents one estimate of discrepancy, under the above judgement. This process is repeated four times for each multi-model member, in order to sample the sensitivity of the optimisation process to the initial point. These difference estimates are then pooled across the multimodel ensemble, giving a sample of four times the number of ensemble members. The mean of these is taken as our estimate of the mean value of discrepancy, and the covariances of the differences about the ensemble mean serve as our estimate of the discrepancy covariance matrix, after allowing for a component due to internal climate variability.

This approach allows us to provide projections combining results from perturbed physics and multi-model ensembles, thus avoiding exclusive reliance on results from the Hadley Centre model. The slab models used in the discrepancy calculation were selected from those contributed to the IPCC AR4 (Randall *et al.* 2007), and the Cloud Feedback Model Intercomparison Project (CFMIP) (e.g.

Webb *et al.* 2006), using data interpolated to the HadSM3 model grid. Some models could not be used as insufficient data was available, and one model was excluded because the design of its simulation of the response to doubled CO₂ excluded the contribution of surface albedo changes from melting sea-ice, this being a process of known importance included in the other models. In the remaining 14 models, data was available for nearly all of the required variables, but with isolated exceptions (mainly daily data required to calculate the required indicators of temperature and precipitation extremes, which was missing from five of the models). Here, values of the missing variables were estimated from inter-variable correlations derived from the multi-model ensemble. In two cases where more than one model was potentially available from a given institute, statistical tests showed that these models could not reasonably be assumed to give quasi-independent estimates of model error, so the model variant thought to be less credible (based on criteria of lower resolution in one case, and published assessments by the relevant modelling centre in the other) was excluded. This left 12 models to be used in the discrepancy calculation (Table 3.1).

Model Name	Modelling Centre	Source
UIUC	University of Illinois, USA	CFMIP
MIROC3.2medres	Centre for Climate System Research, National Institute for Environmental Studies, Frontier Research Centre for Global Change, Japan	CFMIP
MIROC3.2hires	Centre for Climate System Research, National Institute for Environmental Studies, Frontier Research Centre for Global Change, Japan	IPCC
HadGSM1	Met Office Hadley Centre, UK	IPCC
CGCM3.1 T63	Canadian Centre for Climate Modelling and Analysis, Canada	IPCC
CSIRO-MK3.0	Commonwealth Scientific and Industrial Research Organisation, Australia	IPCC
ECHAM5/MPI-OM	Max Planck Institute for Meteorology, Germany	IPCC
GFDL-CM2.0	Geophysical Fluid Dynamics Laboratory, USA	IPCC
GISS-ER	Goddard Institute for Space Studies, USA	IPCC
INM-CM3.0	Institute for Numerical Mathematics, Russia	IPCC
MRI-CGCM2.3.2	Meteorological Research Institute, Japan	IPCC
NCAR-CCSM3.0	National Center for Atmospheric Research, USA	IPCC

Table 3.1: Climate models used in the estimation of structural errors (discrepancy). Randall *et al.* (2007) (Table 8.1 therein) and Webb *et al.* (2006) summarise some basic features of models sourced from IPCC and CFMIP, respectively, and also provide supporting references for papers giving detailed model descriptions. Note that Table 8.1 of Randall *et al.* describes dynamical ocean-atmosphere configurations of the models, from which are derived the mixed layer ocean-atmosphere (*slab*) configurations used here.

Assumptions and limitations

Whilst this method of calculating discrepancy provides an appropriate means of quantifying uncertainties in projected future changes consistent with current climate modelling technology, it is important to recognise caveats associated with the approach. Firstly, it assumes that the structural errors in different models can be taken to be independent. Whilst there is evidence for a degree of independence (for example, model errors in multiyear climate averages reduce significantly when ensembles of different models are averaged together (e.g. Lambert and Boer, 2001; Reichler and Kim, 2008)), there is also evidence that some errors are common to all models (see Annex 3), due to shared limitations such as insufficient resolution or the widespread adoption of an imperfect parameterisation scheme. From this perspective, our estimates of discrepancy can be viewed as a likely lower bound to the true level of uncertainty associated with structural model errors. However, another caveat is that we do not take into account variations in the credibility of different multi-model ensemble members when calculating discrepancy, partly because there is no widely recognized means of quantifying such variations (Randall *et al.* 2007), and partly because such an exercise would introduce an element of double counting in the use of observations in our Bayesian framework. Nevertheless, the assumption of equal credibility carries the risk that models which simulate climate relatively poorly could yield excessively large estimates of discrepancy, thus overestimating the impact of structural errors.

It is clear, therefore, that the sensitivity of our projections to plausible variations in discrepancy is an important test of their robustness (see Annex 2, and further discussion in Section 3.3). In the case of the historical component of discrepancy, such tests can be augmented by diagnostic checks, since the magnitude of biases in our model simulations can be calculated *a posteriori*. We used our emulator to estimate the location in the model parameter space which gives the best simulation of historical climate, and then calculated the squared error between emulated and observed values found in practice, for each of the variables used in our weighting of different model variants (see Section 3.2.9). For each variable, the squared error was then divided by our *a priori* estimate of its expected value, this consisting of the sum of the variances arising from our prior estimate of discrepancy, observational errors, and emulation errors. The average value of these normalised squared errors was found to be ~0.3, indicating that the structural component of model error may be rather smaller than our *a priori* estimates derived from other climate models without reference to the observations. This suggests that the potential risk that the presence of common systematic errors in models might lead us to underestimate historical discrepancy is not realized in practice, at least for the set of historical observables considered. Obviously we cannot perform corresponding diagnostic checks on the discrepancy attached to future variables, and there is no guarantee that an overestimate in historical discrepancy would necessarily imply a corresponding overestimate of future values.

3.2.9 Use of climate variables to estimate discrepancy and weight projections

The calculation of weights for different locations in the HadSM3 parameter space (Section 3.2.7) requires us to compare emulated estimates of historical climate against some set of corresponding observations. In addition, the calculation of discrepancy (Section 3.2.8) requires us to compare emulated estimates of both historical climate and the response to doubled CO₂ against simulated values from multimodel ensemble members. In this sub-section we describe the set of variables upon which these comparisons are based.

Which observations are used to weight UKCP09 projections?

Our choice of potential observational constraints is restricted to historical variables which can be simulated by our ensemble of HadSM3 simulations, or which can be inferred with acceptable accuracy via the timescaling procedure of Sections 3.2.4 and 3.2.6. This precludes, for example, the use of observations of properties relating to sub-surface ocean, sulphur cycle or terrestrial ecosystem processes (e.g. ocean salinity or temperature cross-sections, net primary productivity of the biosphere, etc.) or of coupled ocean–atmosphere modes of variability in which ocean transport plays a role, such as the El Niño–Southern Oscillation. In the main, therefore, we are restricted to the use of spatial fields of multiannual seasonal means of physical variables describing surface and atmospheric characteristics of recent historical climate. We are also restricted by the set of fields available from the multi-model ensemble used to generate our discrepancy estimates (Section 3.2.8). Nevertheless, this still constitutes a substantial subset of the metrics typically used to assess climate simulations (e.g. Taylor, 2001; Reichler and Kim, 2008). Specifically, we use observed latitude–longitude fields of sea surface temperature, land surface air temperature, precipitation, pressure at mean sea level, shortwave and longwave radiation at the top of the atmosphere, shortwave and longwave cloud radiative forcing, total cloud amount, surface fluxes of sensible and latent heat, and latitude–height distributions of zonally averaged atmospheric relative humidity. This amounts to a very large number of variables, given that a single spatial field consists of ~7000 grid box values. However there are significant spatial relationships within each field, and also relationships between different fields, so it is possible (and necessary, for computational reasons) to capture the main variations found in our ensemble simulations of these observables in a smaller number of independent variables, as described in the following sub-section.

In addition, we also include changes in large-scale features of surface temperature patterns observed during the twentieth century as an additional constraint. This is desirable because the ability to replicate historical temperature changes is widely recognised as an important test of the credibility of projected future changes, and has been used as a formal observational constraint in a number of studies (e.g. Allen *et al.* 2000; Stott *et al.* 2006a,b). It is feasible to do this in UKCP09 because our timescaling technique allows us to infer this aspect of time-dependent historical climate change for any given point in parameter space, by using our simple climate model tuned to the relevant climate sensitivity (Section 3.2.4). We therefore include historical changes in four indices identified by Braganza *et al.* (2003), which capture key features of the characteristic response to increasing greenhouse gases found in climate model simulations, these being the global mean, land–ocean and interhemispheric temperature contrasts and the zonally averaged meridional temperature gradient in Northern Hemisphere mid-latitudes. Stott *et al.* (2006a) show that these indices capture most of the information obtained from comprehensive spatiotemporal analyses of the past warming attributable to forcing from greenhouse gases, aerosols and natural forcing agents, and therefore provide an important constraint on future temperature changes at continental to global scales (e.g. Stott *et al.* 2006b; Kettleborough *et al.* 2007). We also account for structural error in our estimates of the Braganza indices, by combining our emulation and timescaling techniques to *predict* the results of estimates derived from multimodel ensemble members, using an approach consistent with that used to calculate other aspects of discrepancy (see Section 3.2.8).

Expressing observational constraints through a limited set of key variables

Our set of observables, whilst incomplete, constitutes a large collection of

variables covering a variety of physical climate characteristics. This should substantially reduce the risk of erroneously assigning a high weight to a location in parameter space which achieves a good fit to observations through a fortuitous compensation of errors. In order to make our calculations tractable, it is necessary to reduce the number of historical multiannual mean climate variables used in the calculation of relative likelihoods for different parts of parameter space. This is done through an eigenvector analysis, identifying a limited set of orthogonal multivariate patterns which explain the main variations in behaviour found within our ensemble of HadSM3 simulations. Fields of values for each climate variable are expressed in dimensionless units prior to the eigenvector analysis, by normalizing values at each location by the globally averaged value of the standard deviation of the relevant variable across the HadSM3 ensemble. The choice of cutoff for the number of retained eigenvectors is determined by a balance between: (i) the need to include a wide range of historical information in order to identify physically and statistically significant variations in the fit to observations found in different parts of parameter space; and (ii) the need to ensure that a reasonable proportion of points in parameter space receive a non-negligible weight, so that robust projections can be obtained by sampling a large but finite sample of points. Statistical tests indicate that six eigenvectors is the appropriate choice (see also Annex 2). The retained eigenvectors explain 66% of the variance found within the HadSM3 ensemble. The projections of emulated multiyear mean climate onto these six eigenvectors, plus the four Braganza *et al.* indices of large scale historical surface temperature trends, form the set of observables from which the weights are calculated.

Observational uncertainties

The specification of uncertainties associated with the verifying historical observations is in principle an important consideration. For the indices of historical surface temperature changes, the estimates are derived from the error estimates supplied by Brohan *et al.* (2006) for the HadCRUT3 dataset. The available observational climatologies for the multiyear mean variables do not possess comprehensive error estimates, so we take the simpler approach of using two alternative verifying datasets for each variable, and randomly generating plausible alternative observed values by interpolating between the two datasets. Improving the specification of observational uncertainties is an issue for future research.

Which climate variables are used to find perturbed physics analogues to multimodel ensemble members?

As explained in Section 3.2.8, we estimate discrepancy by finding locations in the HadSM3 parameter space which produce emulated estimates of climate which best fit results from the simulations of an ensemble of alternative models. The fit is calculated by combining information from simulations of both historical climate and future climate change. The historical information is based on projections onto the six eigenvectors of spatial patterns of time-averaged climate described above. The future climate change information is provided from six multivariate eigenvectors of the simulated response to doubled CO₂. These are obtained from an eigenvector analysis of patterns of change in the ensemble of perturbed physics simulations, based on the same set of variables used to determine eigenvectors of historical climate (see above). The simulated climate changes of multimodel ensemble members are then projected onto these eigenvectors, as are emulated changes from different points in the HadSM3 parameter space, allowing us to add six coefficients of future climate change to the six historical variables used to determine the best perturbed physics analogues to any given multimodel ensemble member.

Although we use only twelve derived variables in this matching process, these encapsulate information from global patterns of historical climate and future change of a range of basic climate variables. This ensures that it would only be possible to find a good overall match (over different variables and regions) if HadSM3 analogues can be found which closely replicate all aspects of the representations of physical processes found in any given multimodel ensemble member. Any outstanding mismatch (beyond the effects of internal climate variability) should then arise from the true effects of structural differences between HadSM3 and the multimodel ensemble member, and can be taken as an estimate of discrepancy.

3.2.10 Probabilistic projections of the equilibrium response to doubled carbon dioxide

As explained in Sections 3.2.7–3.2.9, probabilistic projections of equilibrium climate changes in response to doubled CO₂ provide the cornerstone of the UKCP09 methodology. This process produces projections of changes in the UKCP09 variables at five global climate model (HadSM3) grid boxes covering the UK landmass (and also a further nine points covering surrounding marine regions), for every month of the year. Here we provide a few illustrations of how this part of the method works in practice, and what criteria are considered in assessing the credibility of the results.

Figure 3.4 shows an example, for changes in the 20-yr average of surface air temperature (T_{mean}) over Wales, in March. The green histogram shows our perturbed physics ensemble of 280 HadSM3 simulations, while the multi-model ensemble (MME) results are shown as black ticks along the x-axis. The MME results provide a means of estimating the impact of structural errors in HadSM3, via the discrepancy term described in Section 3.2.8. We estimate discrepancy by taking each MME member in turn, and use a search algorithm to find four locations within the HadSM3 parameter space which match the results of the MME member most closely, based on multivariate global patterns of both historical climate and changes in response to doubled CO₂ (see Section 3.2.9). Once the four HadSM3 analogues have been found, discrepancy values can be calculated

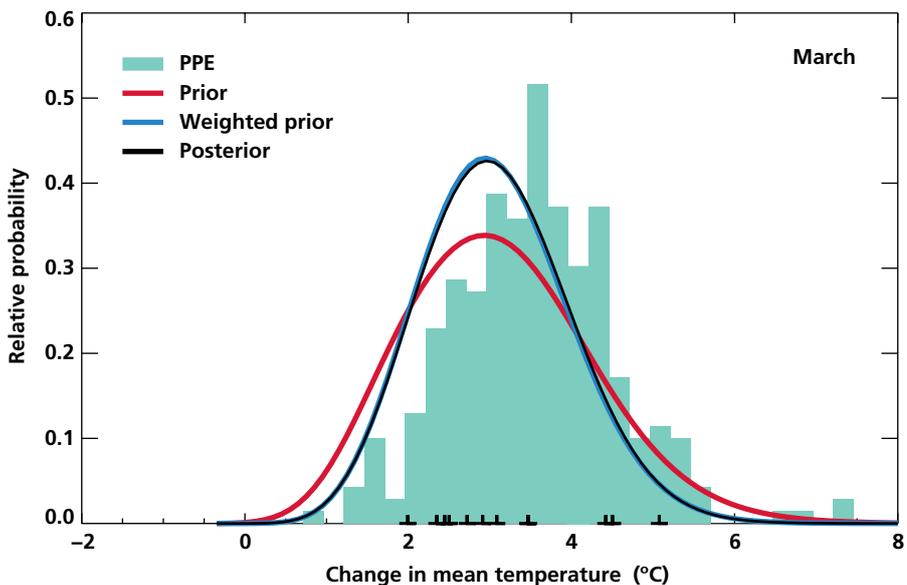


Figure 3.4: Changes in 20 yr-mean surface air temperature (T_{mean}, °C) over the HadSM3 grid box corresponding to Wales, in March, in response to doubled CO₂. Green histogram shows 280 perturbed physics simulations of HadSM3. Black ticks show corresponding changes simulated by 12 multi-model ensemble members. Red curve shows the distribution obtained by emulating responses across the full parameter space of surface and atmospheric processes in HadSM3. The red curve also includes the broadening effect of adding the variance (but not the mean) of discrepancy. The blue curve shows the effects of weighting the emulated responses according to observational constraints (see Section 3.2.9). The black curve shows the posterior distribution, which includes the shift arising from adding in the mean effect of discrepancy.

for any variable of interest (e.g. temperature change over Wales in March). This is done by applying our emulator to estimate projected changes from the four HadSM3 variants, and comparing those with the simulated projection of the corresponding variable from the target MME member. Repeating this procedure for each of the 12 MME members gives 48 discrepancy estimates in total, from which a mean and variance can be calculated (we assume the discrepancy distribution to be Gaussian).

The coloured curves in Figure 3.4 show how we build up our probabilistic projection from the model simulations. We use our emulator trained on the perturbed physics ensemble results (see Section 3.2.3) to estimate results for a much larger ensemble of model variants sampling the full parameter space of HadSM3. This gives us the red curve, which also contains the impact of the variance of discrepancy (but not the mean value of discrepancy, as we wish to illustrate the impact of this separately). In Figure 3.4 the sampling of the full parameter space, combined with the addition of discrepancy variance, leads to a slight broadening of the distribution of possible changes (red curve cf. green histogram). The median value is also shifted slightly towards a smaller warming, this being an effect of the improved sampling of parameter space inherent in the red curve. We also weight points in parameter space according to emulated estimates of the set of historical climate variables described in Section 3.2.9. This weighting process constrains the emulated projections according to the fit to observations, and will in general alter the characteristics of the probability distribution of projected changes. In Figure 3.4 the probabilities of small or large temperature increases are reduced by the weighting (blue curve cf. red curve), while the probabilities of intermediate changes increase somewhat. The mean discrepancy is then added to the projected changes at each location in the HadSM3 parameter space, to produce the final (posterior) probabilistic projection (black curve cf. blue curve).

We cannot make a blanket assumption that this procedure will lead to the production of a credible result. For example, a basic assumption of our approach is that robust probabilities would be difficult to infer from small multi-model ensembles in isolation (see Section 3.1), and that perturbed physics ensembles

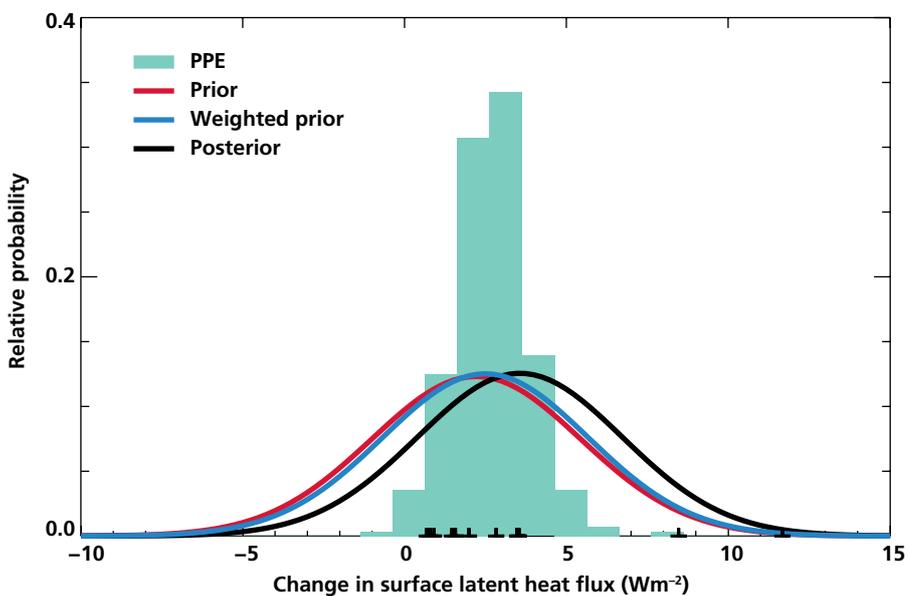


Figure 3.5: As Figure 3.4, for changes in surface latent heat flux (Wm^{-2}) over the HadSM3 grid box corresponding to Scotland, for September–November.

are therefore needed to supply a more systematic means of sampling key process uncertainties to first order. If this is the case, then we would expect the spread of changes simulated by the perturbed physics ensemble to encompass that described by the multi-model ensemble, as it does in Figure 3.4.

We checked all the UKCP09 variables according to this criterion, and generally found that the spread of MME responses did lie within that of the HadSM3 ensemble. For surface latent heat flux, however, two MME members were often found to give projections at or beyond an extreme of the range given by our HadSM3 ensemble (Figure 3.5 shows a typical example). This signals that for latent heat flux the simulated changes are strongly dependent on detailed choices made in the physics of different climate models, and cannot be assumed to be approximately independent of how our experimental design was constructed (for example our decision to base the perturbed physics ensemble on HadCM3/HadSM3, rather than on some other climate model). In Figure 3.5 the outlying MME responses lead to a large discrepancy variance, which substantially inflates the spread in the red, blue and black curves, leading in particular to the projection of a significant probability for negative change in latent heat flux. This is not supported by any of the underlying model simulations. We therefore conclude that the method cannot be used to provide robust probabilistic projections for latent heat flux.

Another issue concerns the magnitude of the shift in the final projections resulting from the mean of the discrepancy term (black cf. blue curve in Figure 3.4). If the perturbed physics ensemble is an effective means of sampling key uncertainties to first order, we would expect the mean value of discrepancy to exert a limited (albeit non-trivial) influence on the final results. This is indeed the case in Figure 3.4. Here, it is important to understand that the mean discrepancy can in theory be large, even when the multi-model and perturbed physics ensemble results cover similar ranges. This is because the procedure used to match MME members to their nearest perturbed physics ensemble analogues is conducted using information based on a wide range of historical and future climate information derived from global multivariate patterns. This is done to ensure that it will only be possible to find a perfect match (across all variables and regions) if the perturbed physics analogues truly replicate all aspects of the representations of physical processes simulated in their target MME members. Any remaining disparities (for some particular local variable like temperature change over Wales in March) will then be a consequence of true structural differences between HadSM3 and the MME members. Note that if we had attempted to calculate the discrepancy by conducting the matching exercise using a more limited choice of variables (say using only temperature changes over the UK), we would have risked finding misleadingly good matches over the chosen variables (through a convenient local compensation of errors effectively achieved via statistical overfitting), accompanied by unrealistically poor matches over other variables or regions not included in the matching process.

Figure 3.6 shows a histogram of the shifts in Tmean arising from the mean of the discrepancy, considering the 60 Tmean projections obtained by pooling monthly changes at all five UK land points in HadSM3. In most cases the mean discrepancy is within the range plus or minus 0.5°C (as in Figure 3.4), and therefore provides a significant but not dominant contribution to the final projection, compared to the spread of responses simulated by the HadSM3 ensemble, or emulated across the full HadSM3 parameter space. In such cases, we typically find that the median of the posterior distribution lies somewhere between the medians of the HadSM3 and MME ensembles.

Occasionally, however, larger shifts are found. Figure 3.7 shows the biggest shift (between the posterior probabilistic projection and the underlying climate model simulations) found in our Tmean projections, over Scotland in March. In this particular case the median of the posterior distribution ends up towards the lower end of the distributions of both the HadSM3 and MME simulations, because all the effects described above (sampling the full parameter space, weighting, and discrepancy) conspire to shift it in the same direction. The largest component in the total shift comes from discrepancy. Detailed investigation reveals that this occurs because the HadSM3 ensemble members have a larger local snow albedo feedback in their response to doubled CO₂, compared to the MME members. This is due to a cold bias in their present day simulations over Scotland, which means that there is too much snow to melt when CO₂ is doubled in their climate change simulations. The discrepancy calculation captures the resulting bias in their simulated changes, reducing the estimated warming to account for the excessive contribution from reduced snow cover in HadSM3. If this was the only contribution to the total shift, then the median of the posterior distribution (black curve) would in this case lie close to the median of the MME results. However the effects of sampling the full HadSM3 parameter space (red curve cf. green histogram in Figure 3.7), and of weighting the projections with observations (blue curve cf. red curve), both add to the total shift, explaining why the posterior distribution shows a median warming smaller than that of either the HadSM3 or MME ensembles. The posterior distribution thus suggests a probability of about 15% for a warming smaller than those simulated by any of the climate model runs. We believe that the shifts arising from sampling parameter space and weighting are both credible, because these aspects of the method improve the sampling of uncertainties and give more emphasis to the better HadSM3 model variants. We also believe the direction of the shift arising from discrepancy is physically credible (see above). Despite this, the magnitude of the shift in this particular case is a cause for concern, as it must be regarded as uncertain (as explained in Section 3.2.8), and yet exerts a substantial influence on the final result. If Figure 3.7 was a typical example of the impact of discrepancy, it would be difficult to justify the production of probabilistic projections of Tmean.

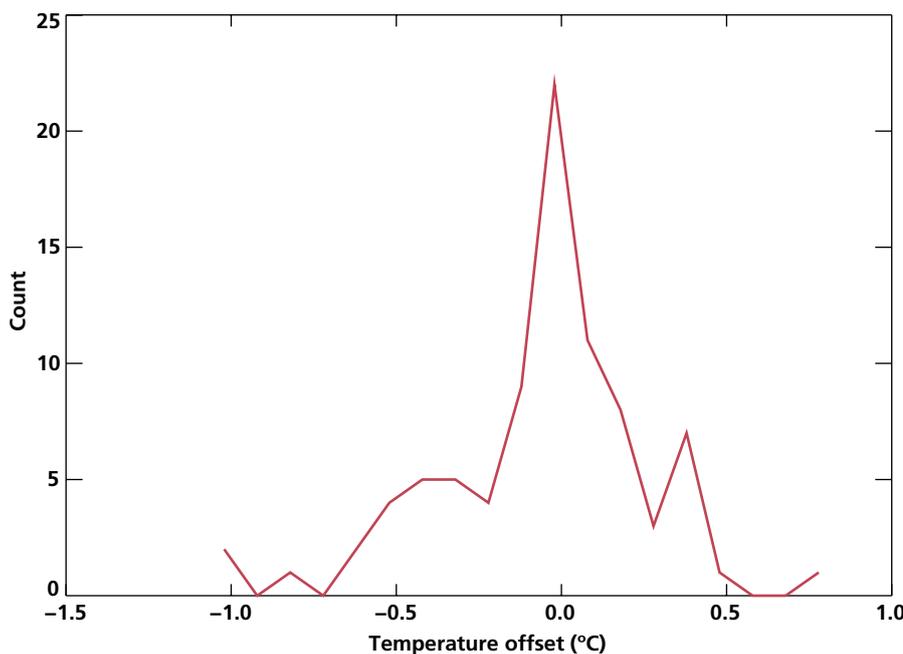


Figure 3.6: Histogram of values for the mean discrepancy for 20 yr mean changes in monthly surface air temperature (°C) in response to doubled CO₂, at UK grid points in HadSM3 (5 grid points x 12 months gives 60 values in all, distributed in bins of width 0.1°C).

However Figure 3.7 is actually an extreme example (see above discussion of Figure 3.6), so overall we judge the impact of discrepancy to be sufficiently modest to justify the production of probabilistic projections for Tmean.

We checked the impact of the shift due to the mean discrepancy in all UKCP09 variables. While isolated examples of significant shifts could be found for some variables (as in Figure 3.7 for Tmean), the typical impacts of such shifts were judged sufficiently modest to imply that the methodology could be considered a reasonable basis for the production of probabilistic projections. However, we note that surface latent heat flux was excluded (due to the mismatch between the MME and HadSM3 ensemble results discussed above). Also, it was not possible to produce probabilistic projections of snowfall or soil moisture content for other reasons, discussed in Section 3.3.

3.2.11 Downscaling for UKCP09

Regional climate model simulations

In order to provide climate projections at the fine spatial scales required for UKCP09 (see Figure 1.2(a), a downscaling method is required to derive such information from our global climate model simulations, run using a horizontal resolution of ~300 km. This was achieved by running simulations of a high resolution limited area regional climate model (RCM), configured from HadCM3 and run at 25 km horizontal resolution. A perturbed physics ensemble of 17 RCM variants was produced, eleven of which were eventually used in UKCP09 (as explained below). These simulations sampled uncertainties in the effects of varying regional physical processes on the simulation of fine scale detail. The simulations capture detailed regional effects of mountains, coastlines and variations in land surface properties, although they do not allow for variations of land surface types within a model grid box, in contrast to a more recent version (Essery *et al.* 2003) being used in additional work to provide a more sophisticated assessment of Urban Heat Island effects (see Annex 7).

Each ensemble member was driven from 1950 to 2100 by time series of lateral boundary conditions (atmospheric surface pressure, wind, temperature and moisture plus chemical species required for the calculation of sulphate aerosol

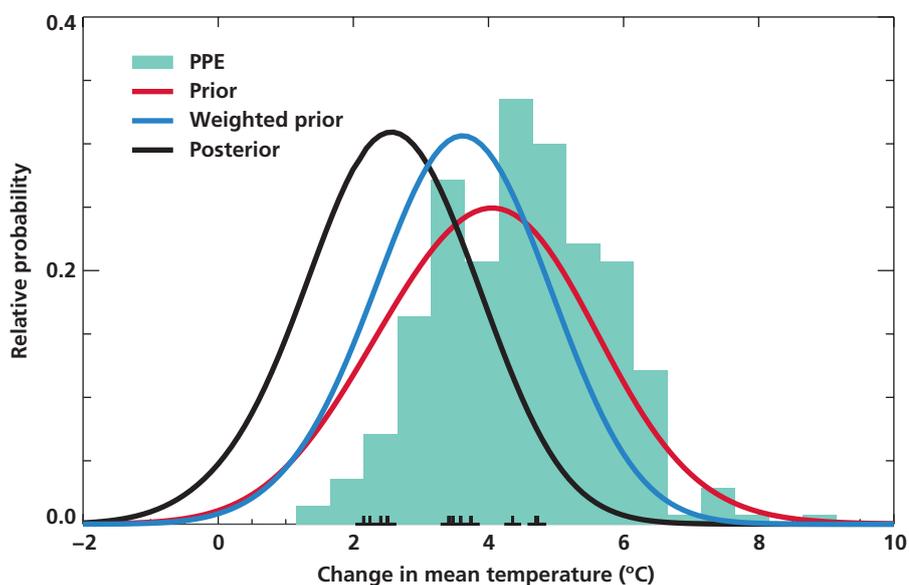


Figure 3.7: As Figure 3.4, for changes in Tmean over the HadSM3 grid box corresponding to Scotland, in March.

concentrations) and surface boundary conditions (sea surface temperatures and sea ice extents) saved from a member of the PPE_A1B ensemble of HadCM3 simulations (Section 3.2.4).^{*} Parameter settings in each RCM ensemble member were chosen to be consistent with the settings used in the relevant HadCM3 simulation. For most parameters this was achieved simply by using the same values in both simulations, however in a few cases the parameters were adjusted to allow for known dependencies on horizontal resolution.

The RCM simulations used the domain shown in Figure 3.8, chosen so as to be large enough to avoid the risk that relaxation to GCM data at the lateral boundaries will damp the simulation of fine scale detail over interior regions of interest (e.g. Jones *et al.* 1995), yet small enough to minimise the risk that inconsistencies could develop between the simulations of large scale climate features in the driving GCM and nested RCM integrations (e.g. Jacob *et al.* 2007).

In eleven ensemble members this experimental design succeeded in producing simulations of detailed climate variability and change over the UK which were physically plausible, and consistent with the driving GCM simulations of

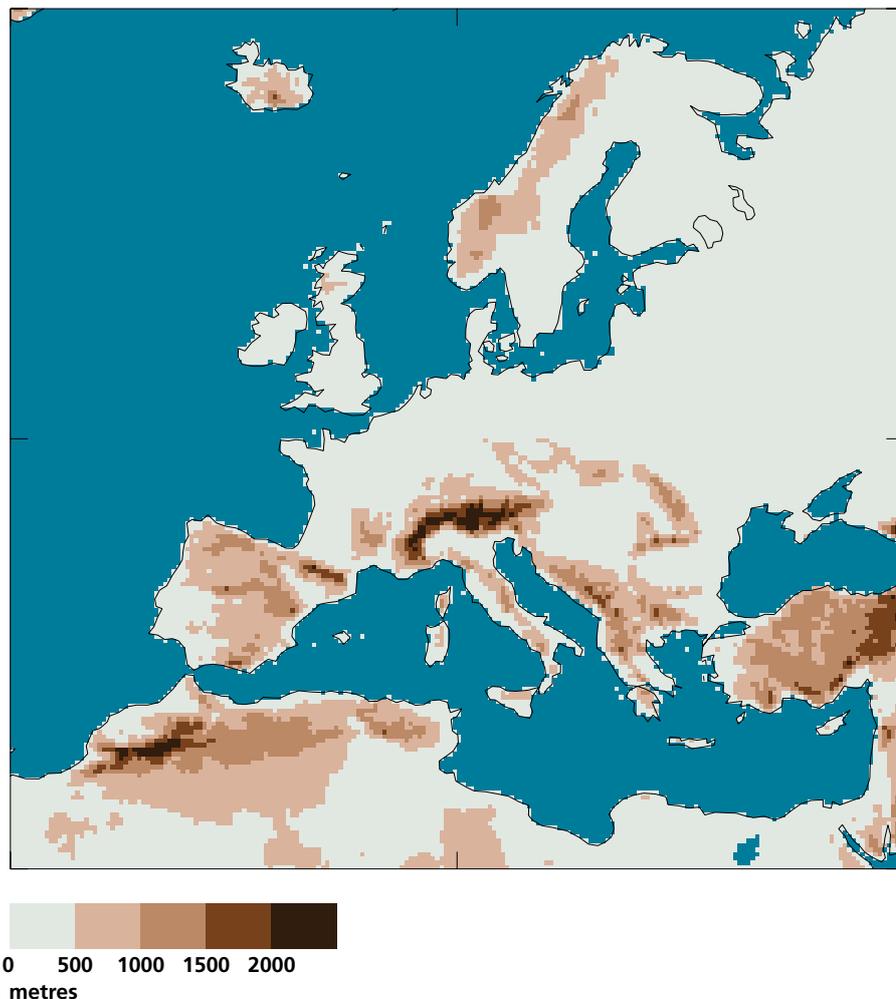


Figure 3.8: Domain used for the UKCP09 regional climate model simulations, excluding the exterior rim within which the model is relaxed to the boundary data supplied from the driving global model simulations. Orographic heights (in metres) are also shown.

^{*} The RCM simulations in UKCP09 are a significant development from those done for UKCIP02 in terms of resolution (25 km cf. 50 km), ensemble design (eleven simulations sampling modelling uncertainties cf. three simulations sampling only initial state uncertainties), and length of simulation (covering 1951–2100 continuously, cf. two *time slices* of 1961–1990 and 2071–2100). These developments allow us to sample a spread of possible realisations of fine scale detail throughout the 21st century in UKCP09, thus avoiding the assumption in UKCIP02 that a single *master pattern* for the 2080s can be scaled back in time to earlier periods.

synoptic scale features (see Annex 3). In six ensemble members, however, the RCM simulations were found to be deficient in their simulations of storms and precipitation, exhibiting too little variability and too many dry days, especially in summer. This was traced to the impact of one of the parameter perturbations, involving a reduction in the order of the diffusive damping applied when calculating dynamical transport of heat, momentum and moisture. The GCM uses sixth order diffusion in its standard variant, whereas the RCM uses fourth order damping as standard (due to its finer grid). Some of our perturbed GCM simulations used fourth order diffusion (thus sampling the effect of increasing the spatial scale of the applied damping), leading to modest reductions in storminess and precipitation variability. An attempt was made to implement an equivalent perturbation in the relevant RCM simulations, moving from fourth to second order diffusion with accompanying changes to the diffusion coefficient to achieve a corresponding change in damping characteristics based on theoretical calculations. However, in practice the changes had a much larger impact than anticipated in the RCM simulations, rendering their time series of winds and precipitation inconsistent with those of the driving GCM runs. These six ensemble members were therefore not used in the calibration of our downscaling procedure, summarised in the following paragraph.

Downscaling to UKCP09 target regions

The downscaling was implemented by developing regression relationships between changes simulated by the RCM over regions for which projections are required by UKCP09 (individual 25 km grid boxes and a set of administrative and river-based regions over land (Figure 1.2), plus a set of marine regions (Figure 1.4)), and changes simulated at nearby grid points in the GCM. This task bears some similarities to a traditional statistical downscaling approach, in which a set of large-scale *predictor* variables is used to obtain values of localized *predictand* variables, using relationships trained on historical observations (e.g. Wilby *et al.* 2004). Such methods assume that historical relationships persist into the future, however such an assumption is avoided in our case, as the relationships are trained using future changes in the predictor and predictand variables simulated by the GCM and RCM, since their purpose is to allow us to infer fine-scale changes for parts of the model parameter space for which no RCM simulation is available.

We expressed the simulated change in a given RCM variable at a given grid point as a univariate linear regression (with slope but no intercept) against the change in the same variable simulated in the GCM at a single nearby grid point. Values for five non-overlapping 30-yr periods (1950–1979, 1980–2009, 2010–2039, 2040–2069, 2070–2099) were expressed as changes relative to the UKCP09 baseline period of 1961–1990, and changes for all five periods and all eleven ensemble members were pooled into a single dataset for the calculation of the regression coefficient (and its associated uncertainty), and the residual unexplained variance. The residual is assumed to be normally distributed with zero mean. Figure 3.9 shows an example, in which the red lines represent the regression relationship, with residual obtained from the scatter of the black crosses about the red lines, which arises from a combination of uncertainty in the relationship between changes in the global and regional models, and also from locally generated internal variability in the RCM runs. This simple approach was used in order to minimise the risk of obtaining unrealistic relationships through overfitting. For non-coastal RCM locations over the mainland UK, the GCM point used in the regression was selected from UK land boxes in HadCM3, selecting the nearest point to the target RCM location unless an adjacent HadCM3 box could be found which explained a significantly greater portion of the variance found

in the RCM response. For marine regions, a similar approach was taken, using predictors chosen from marine HadCM3 boxes nearest or adjacent to the target region. When considering coastal RCM mainland points, or points representing small islands (Channel Islands, Hebrides, Orkney, Shetland, etc.), the predictor variables were selected from surrounding GCM land and sea points, to account for the possibility of a dominant maritime influence on climate at these locations.

Figure 3.9 shows close relationships between the global and regional model changes in winter. Figure 3.10 gives further examples, showing that strong relationships can also be found for summer changes, even for extreme variables subject to considerable internal variability, such as the 99th percentile of daily maximum temperature. Nevertheless, the strengths of the downscaling relationships do depend on which variable, season and region is being considered.

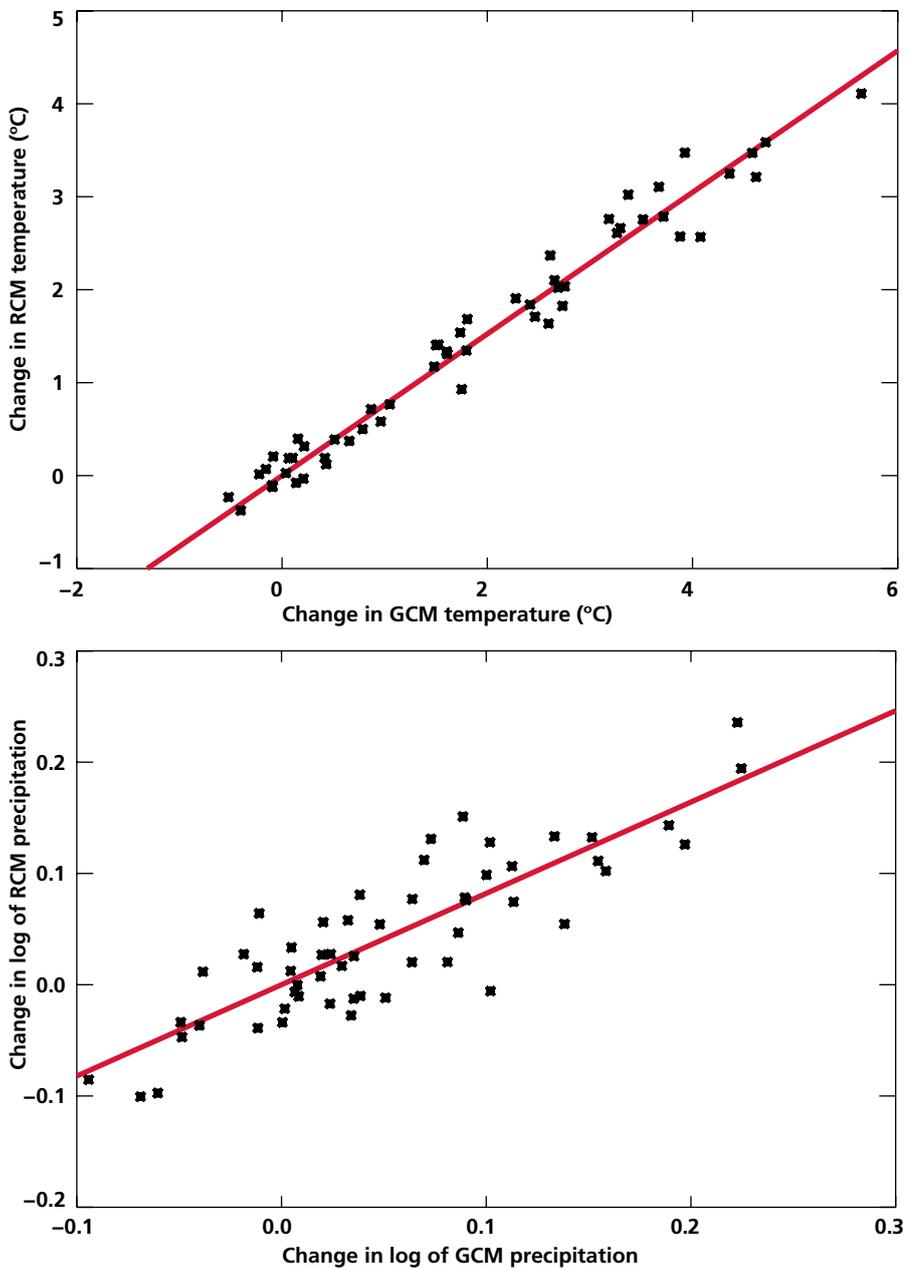


Figure 3.9: Plots of changes in winter surface temperature (°C, top) and in the natural logarithm of precipitation* (bottom), for the North Scotland administrative region, for five non-overlapping 30-yr periods relative to 1961–1990, simulated by 11 members of our regional climate model ensemble (RCM), compared with corresponding changes simulated by driving global climate model simulations (GCM) at a nearby grid point found to be most strongly related to the regional model changes (see text). The red lines show the linear regression relationships between the RCM and GCM changes derived from the data, and used in the downscaling procedure adopted for UKCP09. A zero intercept is imposed on the regression relationships, constraining the red line to pass through the origin and hence preventing the relationship from indicating a non-zero forced response in the RCM when there is no forced response in the GCM.

* Some of the UKCP09 statistical calculations were performed using a transformed variable (here the natural logarithm of precipitation), which is subsequently converted back into the variable provided to users (here percentage changes in precipitation). This is done for reasons explained in Section 3.2.3.

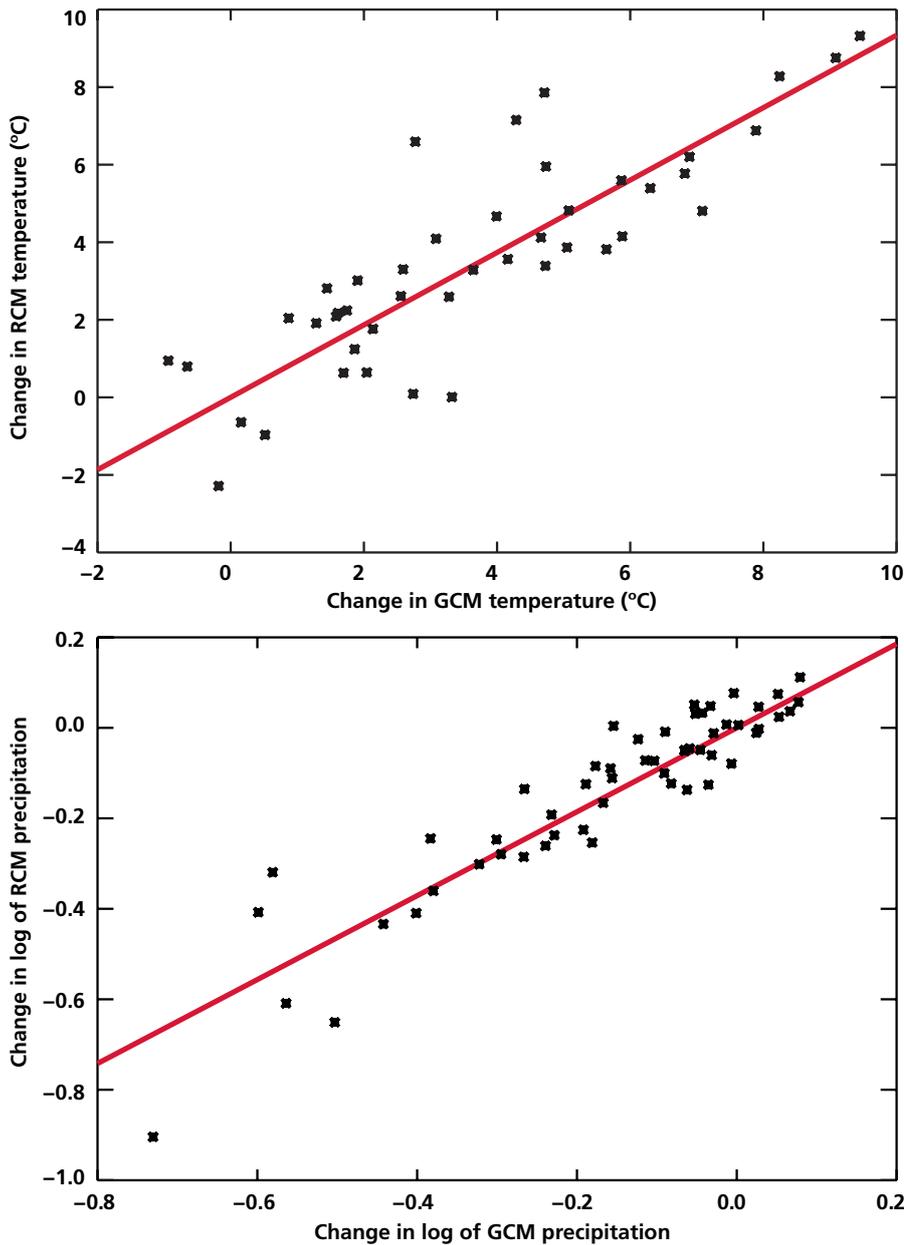
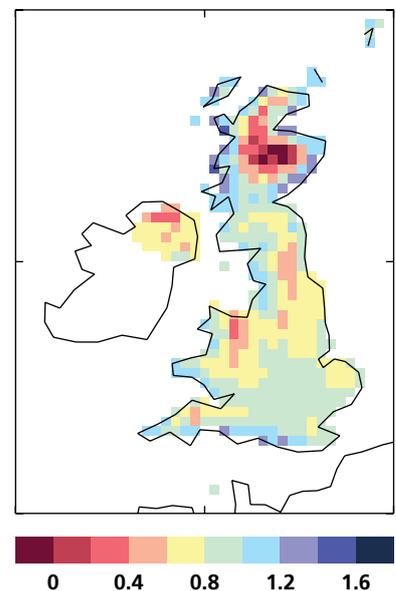


Figure 3.10: As Figure 3.9 for changes in the 99th percentile of daily maximum temperature (°C, top), and in the natural logarithm of precipitation (bottom), for South East England in summer.

Figure 3.11 (below): Plots of regression coefficients between changes in the natural logarithm of winter precipitation in regional and global climate model projections, for UKCP09 25 km grid squares.

Figure 3.11 plots the regression coefficients for changes in winter precipitation at 25 km grid squares around the UK. Significant regional variations are apparent: For example the coefficients exceed unity at many coastal locations, indicating enhanced responses in the RCMs compared with the corresponding GCM simulations, while smaller coefficients are found over parts of Wales, northern England and northern Scotland. Note that the occurrence of small regression coefficients does not necessarily indicate a failure of the downscaling method. For example, this can occur simply because: (i) the RCMs give systematically smaller changes than are found in the GCM simulations, perhaps due to the influence of regional surface topography in modifying changes found at larger scales; or (ii) because the responses in the RCM are dominated by locally generated internal variability. The region of small coefficients over central parts of northern Scotland, for example, occurs because the ratio of internal variability to forced changes is larger than in the driving GCM simulations. However, in some cases our reliance on a simple regression technique using only a single GCM predictor may limit the extent to which the relationship between forced changes in the RCM and GCM simulations is captured in the downscaling procedure.



Assumptions and limitations

Probabilistic projections for UKCP09 target regions were obtained by applying the calibrated downscaling relationships to probabilistic projections of 21st century climate change for the above-mentioned GCM grid boxes covering the UK and surrounding sea points, and hence obtaining estimates for the regions of Figure 1.2 (see Section 3.2.12 for more details). In doing so, a number of limitations of our approach should be recognised. Firstly, we assume that the downscaling relationship (for a given target region and climate variable) is independent of the climate model parameter settings, and of the future period of interest. Secondly, we do not account for variations across parameter space in the skill in simulations of historical fine scale climate features found in our RCM simulations, hence the observational constraints applied to weight different parameter combinations in our Bayesian calculation (see Sections 3.2.7 and 3.2.9) are based purely on aspects of global model performance. Thirdly, we do not account for potential structural errors in our downscaling procedure, arising, for example, from our exclusive reliance on RCM variants configured from HadCM3, or (as noted above) from our neglect of more complex regression techniques based on multivariate GCM predictor variables. All of these limitations arise from the small size of our ensemble of RCM simulations: In particular, we do not possess enough simulations to emulate potential variations in fine scale characteristics of historical or future climate across parameter space. Further research in multivariate downscaling techniques and improvements in computing capacity may allow refined estimates of downscaling uncertainty to be produced in future.

3.2.12 Production of probabilistic projection data for UKCP09

Here we summarise the computational procedure used to generate probabilistic projections for UKCP09 for the SRES A1B scenario, from the elements described in the preceding sub-sections. Figure 3.12 gives a schematic overview of the main elements of the procedure, described in more detail below.

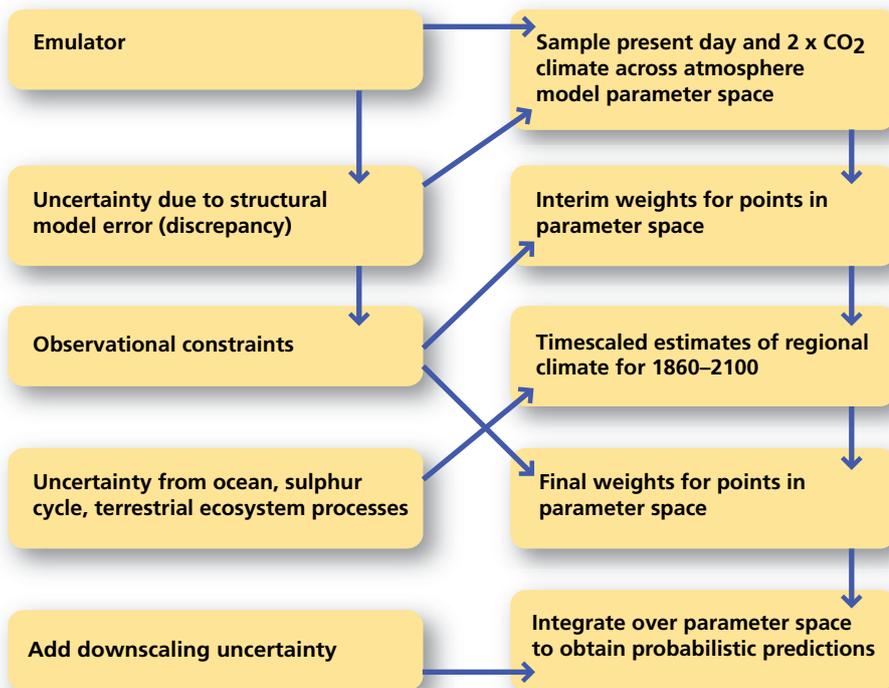


Figure 3.12: Schematic summary of the main elements involved in the derivation of probabilistic projections of climate change for UKCP09, obtained by applying the Bayesian framework of Sections 3.2.7–3.2.9 and the timescaling procedure of Sections 3.2.4 and 3.2.6 to the results of our climate model ensemble simulations. An interim weight, which quantifies the relative likelihood of different model variants based on time-averaged recent climate (see paragraph (i) below), is used to achieve efficient sampling of the atmosphere model parameter space in the timescaling of time-dependent climate changes. Following this final weights are calculated (paragraph (iii)), which account for observations of both recent time-averaged climate, and historical temperature trends.